# Predictability, Complexity, and Learning 

Ilya Nemenman

Co-authored with: William Bialek, Naftali Tishbi

$$
\text { May 8, } 2001
$$

http://xxx.lanl.gov/abs/physics/0007070
http://xxx.lanl.gov/abs/physics/0103076

## Outline

- A curious observation.
- Why a new learning and complexity theory is needed?
- Why and how to use information theory?
- Predictive information, its properties, and relations to other quantities of interest.
- Calculating predictive information for different processes.
- Unique complexity measure through predictive information.
- Possible applications.


## Entropy of words in a spin chain

$$
S(N)=-\sum_{k=0}^{2^{N}-1} P_{N}\left(W_{k}\right) \log _{2} P_{N}\left(W_{k}\right)
$$

For this chain, $P\left(W_{0}\right)=P\left(W_{1}\right)=P\left(W_{3}\right)=P\left(W_{7}\right)=P\left(W_{12}\right)=$ $P\left(W_{14}\right)=2, P\left(W_{8}\right)=P\left(W_{9}\right)=1$, and all other frequencies (probabilities) are zero. Thus, $S(4) \approx 2.95$ bits.

## Entropy of 3 generated chains

- $J_{\mathrm{ij}}=\delta_{\mathrm{i}, \mathrm{j}+1}$
- $J_{\mathrm{ij}}=J_{0} \delta_{\mathrm{i}, \mathrm{j}+1}, J_{0}$ is taken
at random from $\mathcal{N}(0,1)$
every 400000 spins
- $J_{\mathrm{ij}}$ is taken at random
from $\mathcal{N}\left(0, \frac{1}{i-j}\right)$
every 400000 spins
$1 \cdot 10^{9}$ spins total.

Entropy is extensive! It shows no distinction between the cases.

## Subextensive component of the entropy

This component is usually neglected in physics and information theory.

Subextensive entropy shows a qualitative distinction between the cases! What is the significance of this difference?

## Problems in learning and complexity theories

- many frameworks to study learning
- statistical learning theory
- Minimal Description Length (optimal coding of data)
- specific algorithms and learning machines
- psychological and biological analysis of learning and adaptation in animals
- etc.
- different sets of mathematical quantities used
- probabilistic and strict bounds
- learning curves
* in different units (especially, in biology)
- complexities of learning tasks
- etc.
- complexity and (quality) of learning are related-but how?
- many frameworks to study complexity
- Kolmogorov complexity
- Minimal Description Length (stochastic complexity)
- VC-complexity
- causal states (statistical complexity)
- thermodynamic depth
- slow approach of entropy to extensivity (effective measure complexity)
- complexities of dynamical systems
- other entropy-based definitions of complexity
- complexity must be zero for a completely random signal, and some measures get it wrong

There is very little known about connections between various views on learning and complexity.

We need a universal paradigm created, of which all studied problems are special cases.

We base this approach on the notion of predictability.

## Why predictability?

- we learn (estimate parameters, extrapolate, classify, ...) not for the sake of learning; the problem of learning is to generalize and predict from training examples, and estimation of parameters is only an intermediate step
- nonpredictive features in any signal are useless since we observe now and react in the future
- more features to predict is a problem of intuitively higher complexity
- it is impossible to predict a totally random string, so if complexity is based on predictability, for such a string it is zero


## Quantifying predictability

- learning is accrual of information
- Shannon's information theory is the only nonmetric way to quantify information

Thus we will use information theory to study predictability and will define predictive information as the information that the observed data provides about the data that is coming.

## Definitions

| $T, N$ | 0 |
| :---: | :---: |
| past | now future |
| $\mathcal{I}_{\text {pred }}\left(T, T^{\prime}\right)$ | $=\left\langle\log _{2}\left[\frac{P\left(x_{\text {future }} \mid x_{\text {past }}\right)}{P\left(x_{\text {future }}\right)}\right]\right\rangle$ |
|  | $=S(T)+S\left(T^{\prime}\right)-S\left(T+T^{\prime}\right)$ |
| $S(T)$ | $=\mathcal{S}_{0} \cdot T+S_{1}(T)$ |

extensive component cancels in predictive information predictability is a deviation from extensivity!

$$
I_{\mathrm{pred}}(T) \equiv \mathcal{I}_{\mathrm{pred}}(T, \infty)=S_{1}(T)
$$

## Properties of $I_{\text {pred }}(T)$

- $I_{\text {pred }}(T)$ is information, so $I_{\text {pred }}(T) \geq 0$
- $I_{\text {pred }}(T)$ is subextensive, $\lim _{T \rightarrow \infty} \frac{I_{\text {pred }}(T)}{T}=0$
- diminishing returns, $\lim _{T \rightarrow \infty} \frac{I_{\text {pred }}(T)}{S(T)}=0$
- prediction and postdiction are symmetric


## Relations to coding

To code $N+1$ 'st sample after observing $N$ we need, on average, $\ell(N)=-\left\langle\log _{2} P\left(x_{N+1} \mid x_{1}, \cdots, x_{N}\right)\right\rangle=S(N+1)-S(N) \approx \frac{\partial S(N)}{\partial N}$ bits of information.

So we define the universal learning curve that measures excess coding costs due to finiteness of the knowledge we have

$$
\begin{aligned}
\Lambda(N) & \equiv \ell(N)-\ell(\infty) \\
& =S(N+1)-S(N)-\mathcal{S}_{0} \\
& =S_{1}(N+1)-S_{1}(N) \\
& \approx \frac{\partial S_{1}(N)}{\partial N}=\frac{\partial I_{\mathrm{pred}}(N)}{\partial N}
\end{aligned}
$$

## Properties of $\wedge(N)$

- $\lim _{N \rightarrow \infty} \wedge(N)=0$
- integral of $\Lambda(N)$ is the information accumulated about the model
- $\wedge(N)$ relates to conventional learning curves in specific contexts. Example:
- fitting noisy data $\left\{x_{i}, y_{i}\right\}$ with $y=f(x, \boldsymbol{\alpha})$ :

$$
\left\langle\chi^{2}(N)\right\rangle=\frac{1}{\sigma^{2}}\left\langle[y-f(x ; \boldsymbol{\alpha})]^{2}\right\rangle \rightarrow 2 \wedge(N)+1
$$

## Relations to other quantities in learning theory

$\ell(N)$ thermodynamic dive, $N$-th order block entropy, learning curve for some neural networks
$\mathcal{I}_{\text {pred }}(\infty, \infty) \quad$ excess entropy, effective measure complexity, stored information, etc.; tempts to focus on $\mathcal{I}_{\text {pred }}(\infty, \infty)=$ const $<\infty$ - the least interesting cases
$\mathcal{I}_{\text {pred }}(N, \infty)$ analysed as $I(N$, parameters $)$ for parametric models; cumulative information gain, cumulative entropic loss
$I_{\text {pred }}$ universally generalizes all of these quantities!

## How can $I_{\text {pred }}$ behave?

$\lim _{N \rightarrow \infty} I_{\text {pred }}=$ const no long-range structure

- simply predictable (periodic, constant, etc.) processes
- fully stochastic (Markov) processes
$\lim _{N \rightarrow \infty} I_{\text {pred }}=$ panst $\times \log _{2} N \quad$ precise learning of a fixed set of
- learning finite-parameter densities
- analyzed as $I(N$, parameters $)=I_{\text {pred }}(N)$
$\lim _{N \rightarrow \infty} I_{\text {pred }}=\mathrm{const} \times N^{\xi} \quad$ learning more features as $N$ grows
- learning continuous densities
- not well studied


## Problem setup

$Q(x \mid \boldsymbol{\alpha})$ probability density function for $\vec{x}$ parameterized by unknown parameters $\boldsymbol{\alpha}$
$\operatorname{dim} \boldsymbol{\alpha}=K$ dimensionality of $\boldsymbol{\alpha}$, may be infinite
$\mathcal{P}(\boldsymbol{\alpha})$ prior distribution of parameters
$\vec{x}_{1} \cdots \vec{x}_{\mathrm{N}}$ random samples from the distribution

$$
\begin{aligned}
P\left(\vec{x}_{1}, \vec{x}_{2}, \cdots, \vec{x}_{\mathrm{N}} \mid \boldsymbol{\alpha}\right) & =\prod_{\mathrm{i}=1}^{N} Q\left(\vec{x}_{\mathrm{i}} \mid \boldsymbol{\alpha}\right) \\
P\left(\vec{x}_{1}, \vec{x}_{2}, \cdots, \vec{x}_{\mathrm{N}}\right) & =\int d^{K} \alpha \mathcal{P}(\boldsymbol{\alpha}) \prod_{\mathrm{i}=1}^{N} Q\left(\vec{x}_{\mathrm{i}} \mid \boldsymbol{\alpha}\right) \\
S\left(\vec{x}_{1}, \vec{x}_{2}, \cdots, \vec{x}_{\mathrm{N}}\right) \equiv S(N) & =-\int d \vec{x}_{1} \cdots d \vec{x}_{\mathrm{N}} P\left(\left\{\vec{x}_{\mathrm{i}}\right\}\right) \log _{2} P\left(\left\{\vec{x}_{\mathrm{i}}\right\}\right)
\end{aligned}
$$

## Separating the extensive term

$$
\begin{aligned}
S(N)= & -\int d^{K} \overline{\boldsymbol{\alpha}} \mathcal{P}(\overline{\boldsymbol{\alpha}})\left\{d^{N} \vec{x} \prod_{j=1}^{N} Q\left(\vec{x}_{j} \mid \overline{\boldsymbol{\alpha}}\right) \log _{2} \int d^{K} \alpha \mathcal{P}(\boldsymbol{\alpha}) \prod_{i=1}^{N} Q\left(\vec{x}_{i} \mid \boldsymbol{\alpha}\right)\right\} \\
=- & -\int d^{K} \overline{\boldsymbol{\alpha}} \mathcal{P}(\overline{\boldsymbol{\alpha}})\left\{d^{N} \vec{x} \prod_{j=1}^{N} Q\left(\vec{x}_{j} \mid \overline{\boldsymbol{\alpha}}\right)\right. \\
& \times \log _{2} \prod_{j=1}^{N} Q\left(\vec{x}_{j} \mid \overline{\boldsymbol{\alpha}}\right) \int d^{K} \alpha \mathcal{P}(\boldsymbol{\alpha}) \stackrel{\exp \left[-N \mathcal{E}_{N}\left(\boldsymbol{\alpha} ;\left\{\vec{x}_{\mathrm{i}}\right\}\right)\right]}{\left.\prod_{i=1}^{N}\left[\frac{Q\left(\vec{x}_{\mathrm{i}} \mid \boldsymbol{\alpha}\right)}{Q\left(\vec{x}_{\mathrm{x}} \mid \overline{\boldsymbol{\alpha}}\right)}\right]\right\}}
\end{aligned}
$$

This separates $S(N)$ into the extensive and the subextensive terms

$$
\begin{aligned}
\mathcal{S}_{0} & =\int d^{K} \alpha \mathcal{P}(\boldsymbol{\alpha})\left[-\int d^{D} x Q(\vec{x} \mid \boldsymbol{\alpha}) \log _{2} Q(\vec{x} \mid \boldsymbol{\alpha})\right] \\
S_{1}(N) & =-\int d^{K} \bar{\alpha} d^{N} \overrightarrow{x_{i} \mathcal{P}}(\overline{\boldsymbol{\alpha}}) \log _{2}\left[\int d^{K} \alpha P(\boldsymbol{\alpha}) \mathrm{e}^{-N \mathcal{E}_{N}}\right]
\end{aligned}
$$

## Annealed approximation

Under some conditions we may have

$$
\begin{aligned}
\psi\left(\boldsymbol{\alpha}, \overline{\boldsymbol{\alpha}} ;\left\{x_{i}\right\}\right) & \equiv \underbrace{\mathcal{E}_{N}\left(\boldsymbol{\alpha} ;\left\{\vec{x}_{\mathrm{i}}\right\}\right)}_{\text {quenched energy }}-\underbrace{D_{\mathrm{KL}}(\overline{\boldsymbol{\alpha}} \| \boldsymbol{\alpha})}_{\text {annealed energy }} \\
& \equiv-\frac{1}{N} \sum_{\mathrm{i}=1}^{N} \ln \left[\frac{Q\left(\vec{x}_{\vec{x}} \mid \boldsymbol{\alpha}\right)}{Q\left(\vec{x}_{\mathrm{i}} \mid \overline{\boldsymbol{\alpha}}\right)}\right]+\int d \vec{x} Q(\vec{x} \mid \overline{\boldsymbol{\alpha}}) \ln \left[\frac{Q(\vec{x} \mid \boldsymbol{\alpha})}{Q(\vec{x} \mid \overline{\boldsymbol{\alpha}})}\right] \\
& \rightrightarrows 0
\end{aligned}
$$

$$
S_{1}(N) \not \neg S_{1}^{(\mathrm{a})}(N) \equiv-\int d^{K} \bar{\alpha} \mathcal{P}(\overline{\boldsymbol{\alpha}}) \underbrace{\log _{2} \overbrace{\int d^{K} \alpha P(\boldsymbol{\alpha}) \mathrm{e}^{-N D_{\mathrm{KL}}}}^{\text {annealed partition function, } Z(\overline{\boldsymbol{\alpha}} ; N)}}_{\text {annealed free energy, } F(\overline{\boldsymbol{\alpha}} ; N)}
$$

## Density of states

We can rewrite the partition function

$$
\begin{aligned}
Z(\overline{\boldsymbol{\alpha}} ; N) & =\int d D \rho(D ; \overline{\boldsymbol{\alpha}}) \exp [-N D] \\
\rho(D ; \overline{\boldsymbol{\alpha}}) & =\int d^{K} \alpha \mathcal{P}(\boldsymbol{\alpha}) \delta\left[D-D_{\mathrm{KL}}(\overline{\boldsymbol{\alpha}} \| \boldsymbol{\alpha})\right] \\
\int d D \rho(D ; \overline{\boldsymbol{\alpha}}) & =\int d^{K} \alpha \mathcal{P}(\boldsymbol{\alpha})=1
\end{aligned}
$$

The density $\rho$ could be very different for different targets.
Thus learning is annealing at decreasing temperature; properties of predictive information (and learning) almost always depend on $D=0$ behavior of the density.

## Power-law density function

For this case:

$$
\begin{aligned}
\rho(D \rightarrow 0 ; \overline{\boldsymbol{\alpha}}) & \approx A(\overline{\boldsymbol{\alpha}}) D^{(d-2) / 2} \\
S_{1}^{(\mathrm{a})} & \approx \frac{d}{2} \log _{2} N
\end{aligned}
$$

If $d=d(\overline{\boldsymbol{\alpha}})$, then we can get non half-integer coefficients in front of the logarithm term.

- this behavior is known in MDL and other literature
- speed of approach to this asymptotics is rarely investigated


## Examples of the logarithmic predictive information

- Finite parameter models, $\operatorname{dim} \boldsymbol{\alpha}=K$. Then for $\boldsymbol{\alpha} \approx \overline{\boldsymbol{\alpha}}$ and for sound parameterization

$$
\begin{aligned}
D_{\mathrm{KL}}(\overline{\boldsymbol{\alpha}} \| \boldsymbol{\alpha}) & \approx \frac{1}{2} \sum_{\mu \nu}\left(\bar{\alpha}_{\mu}-\alpha_{\mu}\right) \mathcal{F}_{\mu \nu}\left(\bar{\alpha}_{\nu}-\alpha_{\nu}\right)+\cdots \\
\rho(D \rightarrow 0 ; \overline{\boldsymbol{\alpha}}) & \approx \mathcal{P}(\overline{\boldsymbol{\alpha}}) \frac{2 \pi^{K / 2}}{\Gamma(K / 2)}(\operatorname{det} \mathcal{F})^{-1 / 2} D^{(K-2) / 2} \\
\mathcal{F} & - \text { Fisher information matrix }
\end{aligned}
$$

To avoid complications with soundness, we can define the phase space dimensionality of the model family through the exponent of the density function.

- Finite parameter Markov process, Iearn $Q\left(\vec{x}_{1} \cdots \vec{x}_{\mathrm{N}} \mid \boldsymbol{\alpha}\right)$. If energy is extensive,

$$
D_{\mathrm{KL}}\left[Q\left(\left\{\vec{x}_{\mathrm{i}}\right\} \mid \overline{\boldsymbol{\alpha}}\right) \| Q\left(\left\{\vec{x}_{\mathrm{i}}\right\} \mid \boldsymbol{\alpha}\right)\right] \rightarrow N \mathcal{D}_{\mathrm{KL}}(\overline{\boldsymbol{\alpha}} \| \boldsymbol{\alpha})+o(N)
$$

and extensive term is replaced by

$$
\begin{aligned}
S\left[\left\{\vec{x}_{\mathrm{i}}\right\} \mid \boldsymbol{\alpha}\right] & \equiv-\int d^{N} \vec{x} Q\left(\left\{\vec{x}_{\mathrm{i}}\right\} \mid \boldsymbol{\alpha}\right) \log _{2} Q\left(\left\{\vec{x}_{\mathrm{i}}\right\} \mid \boldsymbol{\alpha}\right) \\
& \rightarrow N \mathcal{S}_{0}+\mathcal{S}_{0}^{*} ; \quad \mathcal{S}_{0}^{*}=\frac{K^{\prime}}{2} \log _{2} N
\end{aligned}
$$

then

$$
S_{1}^{(\mathrm{a})}(N)=\frac{K+K^{\prime}}{2} \log _{2} N
$$

Predictive information does not distinguish predictability coming from unknown parameters and from intrinsic long-range correlations. This is similar to describing physical systems with correlations using order parameters.

## Essential singularity in the density function

As $d \rightarrow \infty$ we may imagine the following behavior

$$
\begin{aligned}
\rho(D \rightarrow 0 ; \overline{\boldsymbol{\alpha}}) & \approx A(\overline{\boldsymbol{\alpha}}) \exp \left[-\frac{B(\overline{\boldsymbol{\alpha}})}{D^{\mu}}\right], \quad \mu>0 \\
C(\overline{\boldsymbol{\alpha}}) & =[B(\overline{\boldsymbol{\alpha}})]^{1 /(\mu+1)}\left(\frac{1}{\mu^{\mu /(\mu+1)}}+\mu^{1 /(\mu+1)}\right) \\
S_{1}^{(a)}(N) & \approx \frac{1}{\ln 2}\langle C(\overline{\boldsymbol{\alpha}})\rangle_{\overline{\boldsymbol{\alpha}}} N^{\mu /(\mu+1)}
\end{aligned}
$$

- finite parameter model with increasing number of parameters $K \sim N^{\mu /(\mu+1)} ; S_{1}(N) \sim N^{\mu / \mu+1}, \operatorname{not} S_{1}(N) \sim \frac{N^{\mu / \mu+1}}{2} \log N$
- as $\mu \rightarrow \infty$ complexity grows and then vanishes to the leading order when $S_{1}^{(a)}$ becomes extensive


## Example of the power-law $I_{\text {pred }}$

Learning a nonparametric (infinite parameter) density $Q(x)=$ $1 / l_{0} \mathrm{e}^{-\phi(x)}, x \in[0, L]$, with some smoothness constraints (Bialek, Callan, and Strong 1996).

$$
\begin{aligned}
\mathcal{P}[\phi(x)] & =\frac{1}{\mathcal{Z}} \exp \left[-\frac{l}{2} \int d x\left(\frac{\partial \phi}{\partial x}\right)^{2}\right] \delta\left[\frac{1}{l_{0}} \int d x \mathrm{e}^{-\phi(x)}-1\right] \\
\rho(D \rightarrow 0 ; \bar{\phi}) & =A[\bar{\phi}(x)] D^{-3 / 2} \exp \left(-\frac{B[\bar{\phi}(x)]}{D}\right) \\
S_{1}^{(\mathrm{a})}(N) & \approx \frac{1}{2 \ln 2} \sqrt{N}\left(\frac{L}{l}\right)^{1 / 2}
\end{aligned}
$$

- increasing number of 'effective parameters' (bins) of adaptive size $\sim \sqrt{l / N Q(x)}$
- heuristic arguments for the dimensionality $\zeta$ and the smoothness exponent $\eta$ give $S_{1}(N) \sim N^{\zeta / 2 \eta}$ - demonstrates a crossover from complexity to randomness


## A note on fluctuations

- fluctuations always decrease $S_{1}$
- fluctuations (and $S_{1}$ ) are ill or well defined together with $\mathcal{S}_{0}$
- for finite parameter system fluctuation do not grow with $N$
- for finite Vapnik-Chervonenkis (VC) dimension (capacity measure) fluctuations are uniformly small
- for infinite VC dimension the decrease of fluctuations is prior dependent, very different approaches to asymptotia (even, possibly, phase transitions) are possible
- for infinite parameter systems, fluctuations are necessarily prior (regularization) dependent and are small if sublinear $S_{1}^{(a)}(N)$ can be calculated

Explicit links between statistical learning theory (capacity of model space) and MDL-type theories (volumes in model space) are established.

## Which complexity we study?

We study complexity of predicting a time series, not computational complexity, algorithmic complexity, or similar. So we look for a definition that can be used for

- Occam-style punishment for complexity in statistical inference (statistics)
- defining and measuring complexity of dynamical processes that generate the time series (physics)


## What do we want in complexity measure?

- it must be zero for totally random and for easily predictable processes (accepted among physicists, but not so much among statisticians)
- to relate to physics, it must be measured by conventional thermodynamic quantities (accepted among physicists, but new to statisticians)
- must not be over-universal, that is it should depend not only on entropy (in principle, accepted by everybody, but usually violated by physicists)
- must be an ensemble property (this is controversial, but see Grassberger)
- must relate to specific complexity measures studied before


## Unique measure of complexity!

Complexity measure must be:

- some kind of entropy (we proclaim Shannon's postulates)
- monotonic in $N$ for $N$ equally likely signals
- additive for statistically independent signals
- a weighted sum of measure at branching points if measuring a leaf on a tree
- reparameterization, quantization invariant, thus subextensive
- invertible temporally local transformations (e. g., $x_{k} \rightarrow x_{k}+$ $\xi x_{k-1}$-measuring device with inertia) and prior insensitive *

The divergent subextensive term measures complexity uniquely!
*The last two conditions may be replaced by a requirement that complexity must stay invariant for any choice of the reference distribution (constructed of local operators) that is needed to define entropy of continuous variables.

## What's next?

- separating predictive information from non-predictive using the 'relevant information' technique
- reflection to physics - finding order parameters for phase transitions using behavior of the predictive information
- reflection to biology - large expansion from receptors to primary sensory cortices may be due to efficient representation of predictive information, not current state of the world
- reflection to psychology - experiments on learning distributions and language (power law complexity class) by humans; what expectations of the world do we have?
- reflection to statistics
- nonparametric models may be simpler then finite parameter ones (relevant to biology)
- predictive information is the property of the data (nonparametric extension of the MDL principle)


## Summary

We have built a generalizing and unique theory of learning and complexity.

