Research Overview (2000 – 2010)

Marcus Hutter
Canberra, ACT, 0200, Australia
http://www.hutter1.net/

IDSIA ANU NICTA

Abstract

I will give an overview of what I regard as the core foundations of intelligent systems and of my own research in this direction. I start with a survey of the key subfields and dichotomies. Most, if not all problems can formally be reduced to predicting sequences or acting well. Occam's razor, Bayes' rule, utility, information, and computability theory are key to solving these problems. I will discuss the following developments and applications: model selection based on loss rank, determining (in)dependence of samples by Mutual Information, non-parametric Bayesian inference, robust estimation, Bayesian change point detection, Bayesian sequence prediction, the MDL principle, prediction with expert advice, applications of algorithmic information theory, optimization. computer vision, and image processing. Finally, I briefly summarize my past work in particle physics, medical software development, and others.

Table of Contents

Personal Motivation

1) Work towards a truly and universal Artificial Intelligence system.
$\Longrightarrow$ BSc & MSc & Habil in Theoretical Computer Science.

2) Work on / find the physical Theory of Everything (SuperStrings?)
$\Longrightarrow$ BSc & PhD in Theoretical Particle Physics.

3) Others: P=NP, computer graphics, mathematical and physical puzzles, numerical algorithms, philosophy of science.

My current research is centered around what I believe to be relevant for making progress in (1), namely ...

Primary Research Focus

- Personal goal: Solve the Artificial Intelligence problem.
- Primary research focus: Universal Artificial Intelligence & Information-Theoretic Foundations of Reinforcement Learning
- Importance: Intelligent software already began to and will transform society more than the industrial revolution did two centuries ago.
- Past achievements (2000–2010):
  I've developed the first sound and complete theory of optimal rational intelligent agents, *which spawned*
- many of my other more mundane contributions to information theory, machine learning, and statistics
- Future directions (2011–†):
  Work out the theory (10 year plan)
  Solve the AI problem in practice (25 year plan)

Relevant Research Fields

(Universal) Artificial Intelligence has interconnections with (draws from and contributes to) many research fields:

- computer science (artificial intelligence, machine learning),
- engineering (information theory, adaptive control),
- economics (rational agents, game theory),
- mathematics (statistics, probability),
- psychology (behaviorism, motivation, incentives),
- philosophy (reasoning, induction, knowledge).

Relation between ML & RL & (U)AI

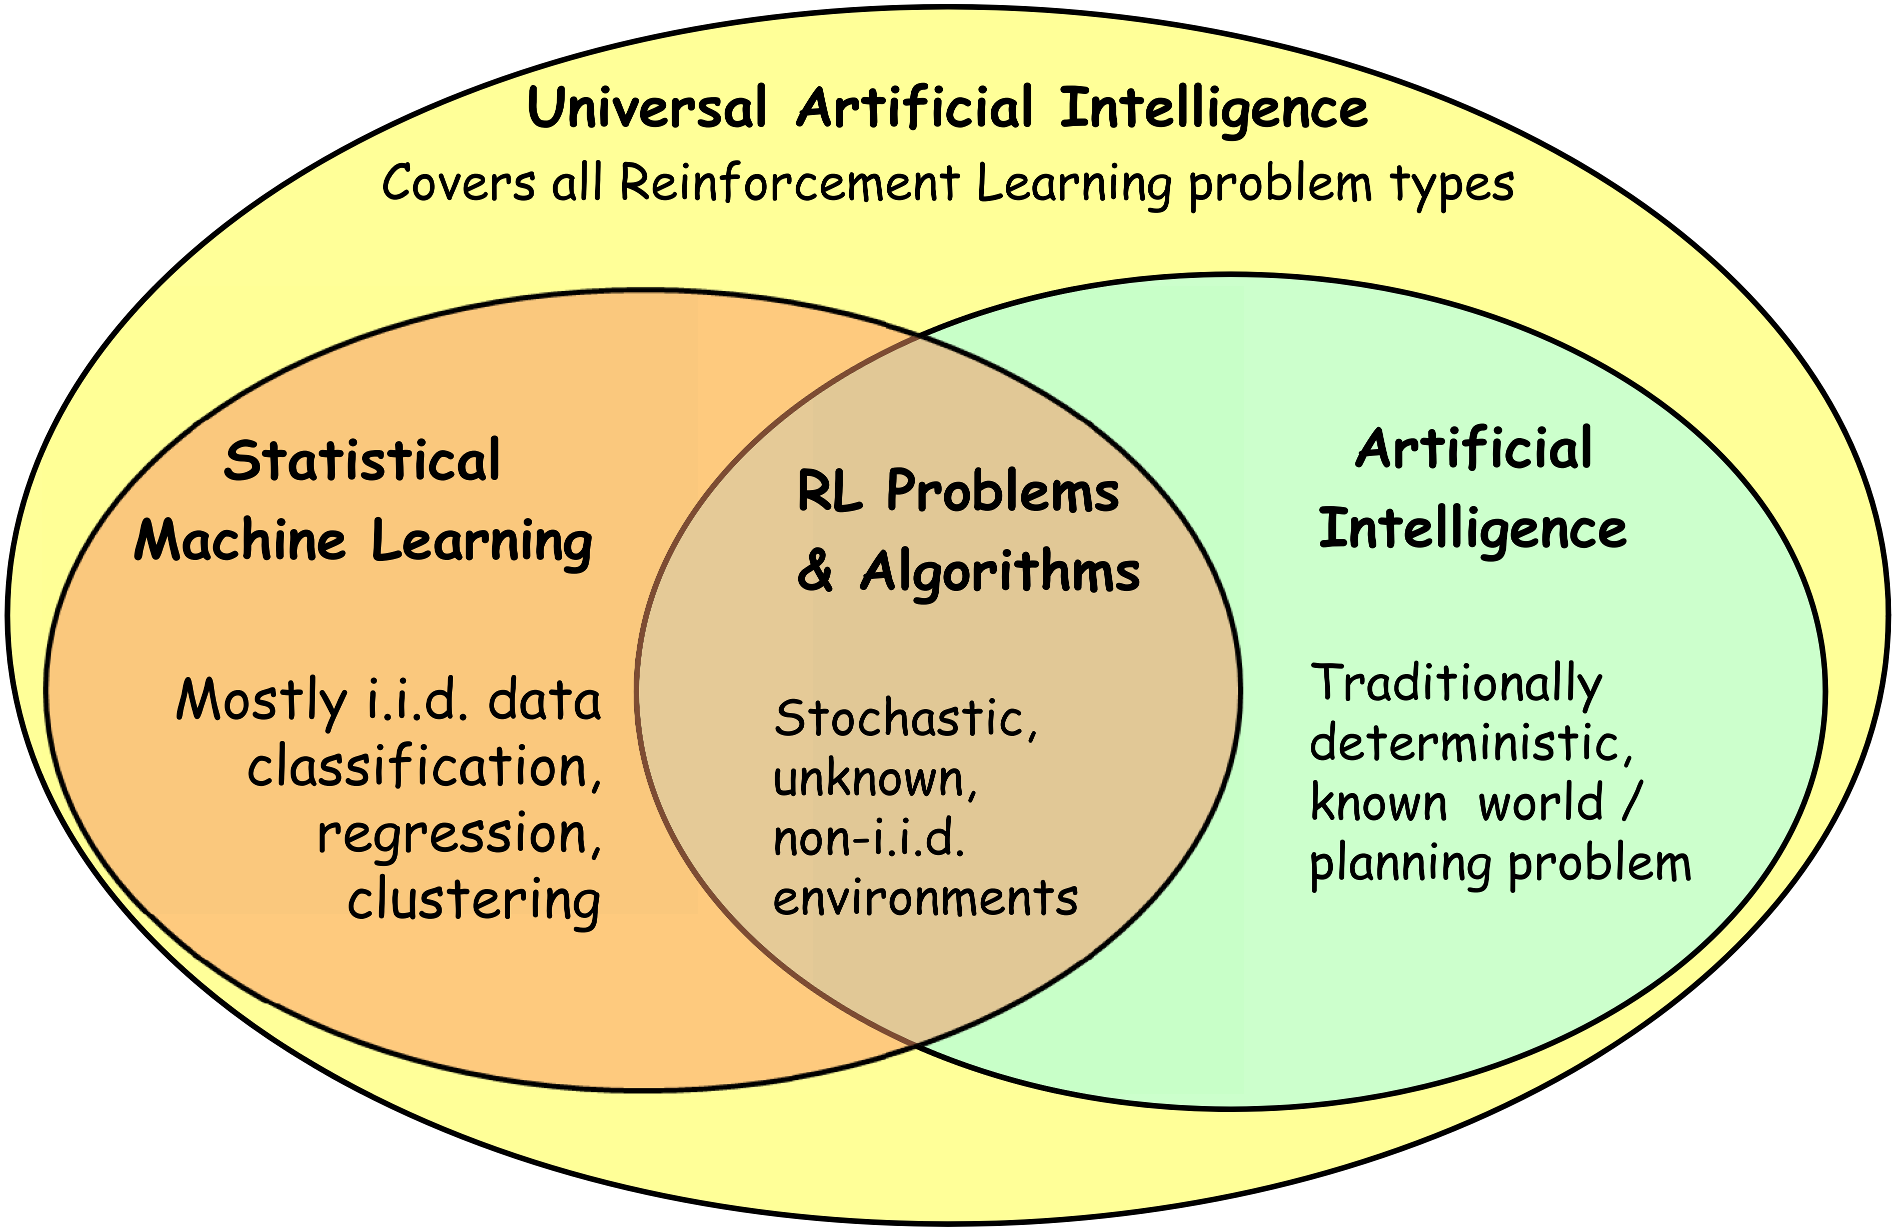

Some Applications of RL

Checkers (Samuel, 1959)
- first use of RL in an interesting real game

(Inverted) helicopter flight (Ng et al. 2004)
- better than any human

Computer Go (UCT, 2006)
- finally some programs with reasonable play

Robocup Soccer Teams (Stone & Veloso, Reidmiller et al.)
- World's best player of simulated soccer, 1999; Runner-up 2000

Inventory Management (Van Roy, Bertsekas, Lee & Tsitsiklis)
- 10-15% improvement over industry standard methods

Dynamic Channel Assignment (Singh & Bertsekas, Nie & Haykin)
- World's best assigner of radio channels to mobile telephone calls

Elevator Control (Crites & Barto)
- (Probably) world's best down-peak elevator controller

Many Robots
- navigation, bi-pedal walking, grasping, switching between skills

TD-Gammon and Jellyfish (Tesauro, Dahl)
- World's best backgammon player. Grandmaster level

My Research Projects and Fields

- machine (reinforcement) learning,
- information theory and statistics,
- Kolmogorov complexity,
- Bayesian/expert/online/sequence prediction,
- minimal description length (MDL) principle,
- computational complexity theory,
- universal Solomonoff induction,
- universal Levin search,
- sequential decision theory,
- adaptive control theory.

See my UAI book on how they fit together →

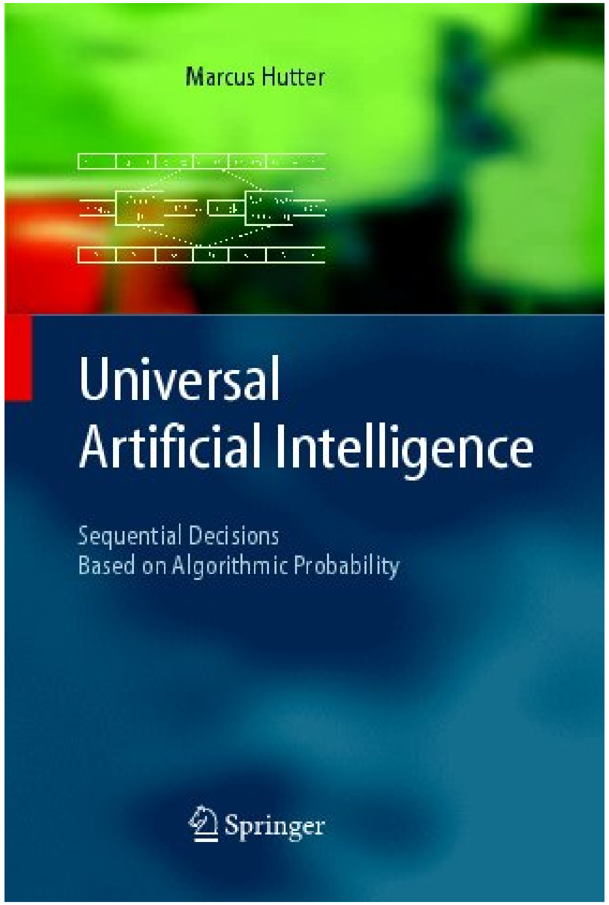

Dichotomies in Intelligent Systems

blue = scope of my research	$\Leftrightarrow$	black = outside my area
rationality	$\Leftrightarrow$	humanness
acting	$\Leftrightarrow$	thinking
(machine) learning / statistical	$\Leftrightarrow$	logic/knowledge-based
online$\approx$lifelong learning	$\Leftrightarrow$	offline/batch/pre-learning
induction $\Rightarrow$ prediction	$\Leftrightarrow$	decision $\Rightarrow$ action
sequential / non-iid	$\Leftrightarrow$	indep. identically distr.
passive prediction	$\Leftrightarrow$	active learning
Bayes $\Leftrightarrow$ MDL	$\Leftrightarrow$	Expert $\Leftrightarrow$ Frequentist
uninformed / universal	$\Leftrightarrow$	informed / problem-specific
conceptual/mathematical	$\Leftrightarrow$	computational issues
exact/principled	$\Leftrightarrow$	heuristic
supervised $\Leftrightarrow$ RL learning	$\Leftrightarrow$	unsupervised learning
learning$\approx$exploration	$\Leftrightarrow$	exploitation$\approx$planning

Sequence Prediction (Bayes/MDL/Expert)

- Setup: Given (non)iid data $D = (x_1, ..., x_n)$, predict x_{n+1}
- Ultimate goal is to maximize profit or minimize loss
- Consider sampling Models/Hypothesis $H_i \in \mathcal{M}$ for D
- Max.Likelihood: $H_{best} = \arg\max_i p(D|H_i)$ (overfits if $\mathcal{M}$ large)
- Bayes: Posterior probability of H_i is $p(H_i|D) \propto p(D|H_i)p(H_i)$
- MDL: $H_{best} = \arg\min_{H_i}\{\text{CodeLength}(D|H_i) + \text{CodeLength}(H_i)\}$
  (Complexity penalization)
- Bayes needs prior(H_i), MDL needs CodeLength(H_i)
- Occam+Epicurus: High prior for simple models with short codes.
- Kolmogorov/Solomonoff: Quantification of simplicity/complexity
- MDL & Bayes work if D is sampled from $H_{true} \in \mathcal{M}$. [H'01-PH'09]
- Prediction with Expert Advice works w/o assumption on D. [HP'05]

Foundations of Induction

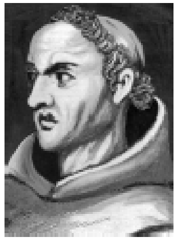

Ockhams' razor (simplicity) principle
Entities should not be multiplied beyond necessity.
Solves e.g. 111111?, 123456?, 141592?,

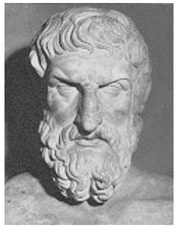

Epicurus' principle of multiple explanations
If more than one theory is consistent with the observations, keep all theories.

Bayes' rule for conditional probabilities
Given the prior belief/probability one can predict all future probabilities.

Turing's universal machine
Everything computable by a human using a fixed procedure can also be computed by a (universal) Turing machine.

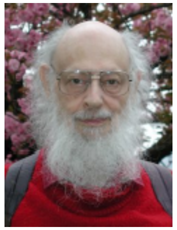

Solomonoff's universal prior=Ockham+Epicurus+Bayes+Turing
Solves the question of how to choose the prior if nothing is known.
$\Rightarrow$ universal induction, formal Occam, AIT,MML,MDL,SRM,...

Human Knowledge Compression Prize

- compression $=$ finding regularities $\Rightarrow$ prediction $\approx$ intelligence
  [hard file size numbers] [slippery concept]
- Many researchers analyze data and find compact models.
- Compressors beating the current compressors need to be smart(er).
- "universal" corpus of data $\Rightarrow$ "universally" smart compressors.
- Wikipedia seems a good snapshot of the Human World Knowledge.
- The ultimate compressor of Wikipedia will "understand" all human knowledge, i.e. be really smart.
- Contest: Compress Wikipedia better than the current record.
- Prize: 50'000 Euro $\times$ the relative improvement to previous record.

[http://prize.hutter1.net]

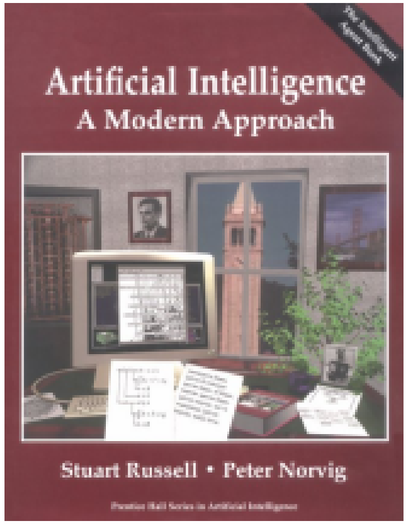

Agent Model for ReActive Problems

Framework for all AI problems!

Is there a universal solution?

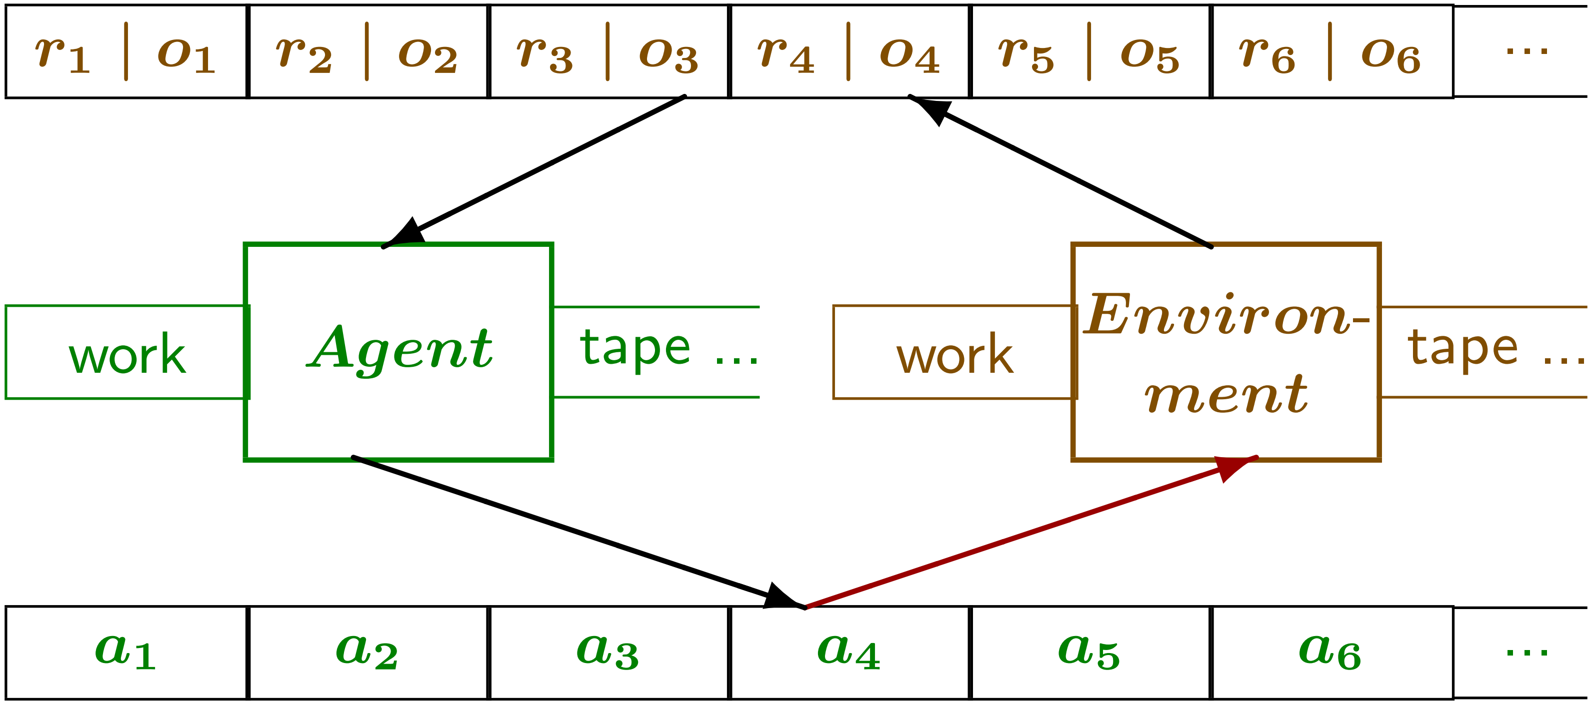

Universal Artificial Intelligence

Key idea: Optimal action/plan/policy based on the simplest world model consistent with history. Formally ...

$$\text{AIXI:} \quad a_k \;:=\; \arg\max_{a_k} \sum_{o_k r_k} \ldots \max_{a_m} \sum_{o_m r_m} [r_k + \ldots + r_m] \sum_{p\,:\,U(p,a_1..a_m)=o_1 r_1..o_m r_m} 2^{-length(p)}$$

k=now, $a$ction, $o$bservation, $r$eward, Universal TM, $p$rogram, m=lifespan

AIXI is an elegant, complete, essentially unique, and limit-computable mathematical theory of AI.

Claim: AIXI is the most intelligent environmental independent, i.e. universally optimal, agent possible.

Proof: For formalizations, quantifications, proofs see $\Rightarrow$

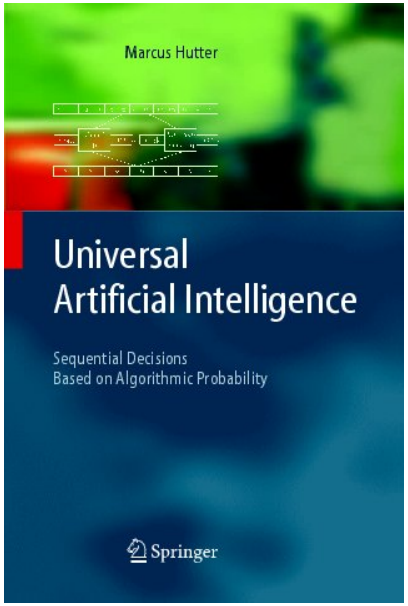

Problem: Computationally intractable.

Achievement: Well-defines AI. Gold standard to aim at. Inspired practical algorithms. Cf. infeasible exact minimax. [H'00-05]

A Monte-Carlo AIXI Approximation

Consider class of Variable-Order Markov Decision Processes.

The Context Tree Weighting (CTW) algorithm can efficiently mix (exactly in essentially linear time) all prediction suffix trees.

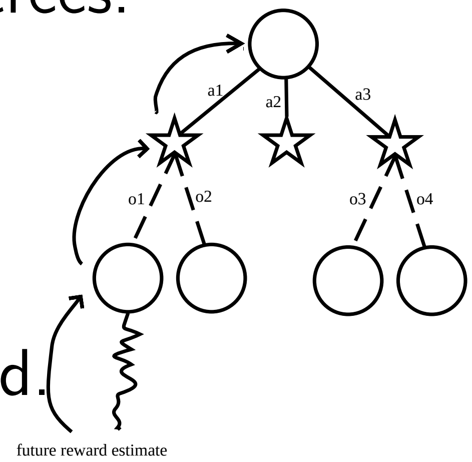

Monte-Carlo approximation of expectimax tree:

Upper Confidence Tree (UCT) algorithm:

- Sample observations from CTW distribution.
- Select actions with highest upper confidence bound.
- Expand tree by one leaf node (per trajectory).
- Simulate from leaf node further down using (fixed) playout policy.
- Propagate back the value estimates for each node.

Repeat until timeout. [VNHUS'09-11]

Guaranteed to converge to exact value.

Extension: Predicate CTW not based on raw obs. but features thereof.

Monte-Carlo AIXI Applications

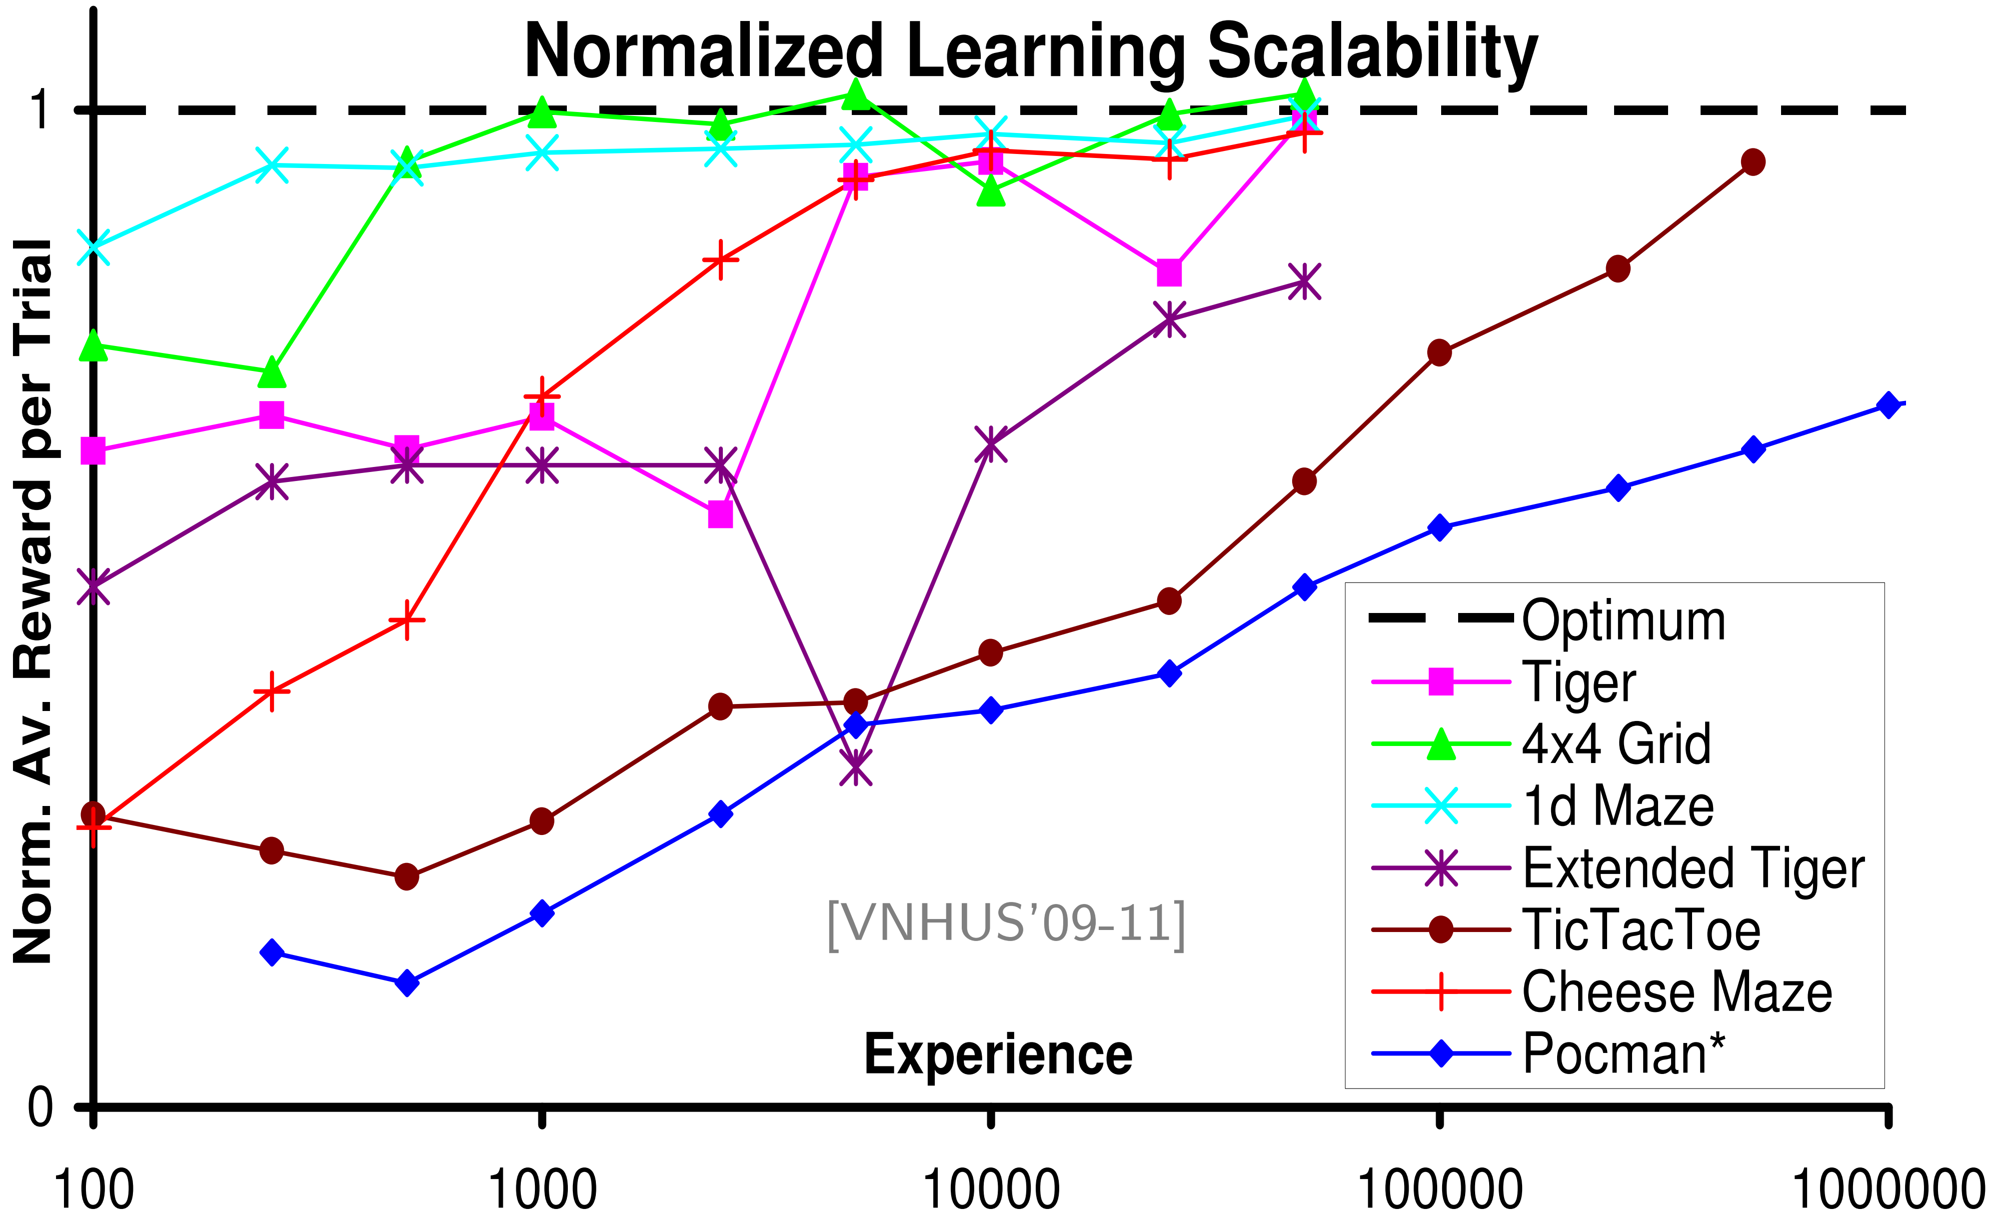

Emphasis in AI/ML/RL $\Leftrightarrow$ Control Theory

Both fields start from Bellman-equations and aim at agents/controllers that behave optimally and are adaptive, but differ in terminology and emphasis:

agent	$\widehat{=}$	controller
environment	$\widehat{=}$	system / plant
(instantaneous) reward	$\widehat{=}$	(immediate) cost
model learning	$\widehat{=}$	system identification
reinforcement learning	$\widehat{=}$	adaptive control
exploration$\leftrightarrow$exploitation problem	$\widehat{=}$	estimation$\leftrightarrow$control problem
qualitative solution	$\Leftrightarrow$	high precision
complex environment	$\Leftrightarrow$	simple (linear) machine
temporal difference	$\Leftrightarrow$	Kalman filtering / Riccati eq.

AIXI is the first non-heuristic formal approach that is general enough to cover both fields. [H'05]

Practical Universal AI/RL (ΦMDP)

Goal: Develop efficient general purpose intelligent agent.

State-of-the-art: (a) AIXI: Incomputable theoretical solution.
(b) MDP: Efficient limited problem class.
(c) POMDP: Notoriously difficult. (d) PSRs: Underdeveloped.

Idea: ΦMDP reduces real problem to MDP automatically by learning.

Accomplishments so far: (i) Criterion for evaluating quality of reduction.
(ii) Integration of the various parts into one learning algorithm.
(iii) Generalization to structured MDPs (DBNs). [H'09]
(iv) Theoretical and experimental investigation. [SH'10,NSH'11]

ΦMDP is promising path towards the grand goal & alternative to (a)-(d)

Problem: Find reduction Φ efficiently (generic optimization problem?)

Intelligent Agents in Perspective

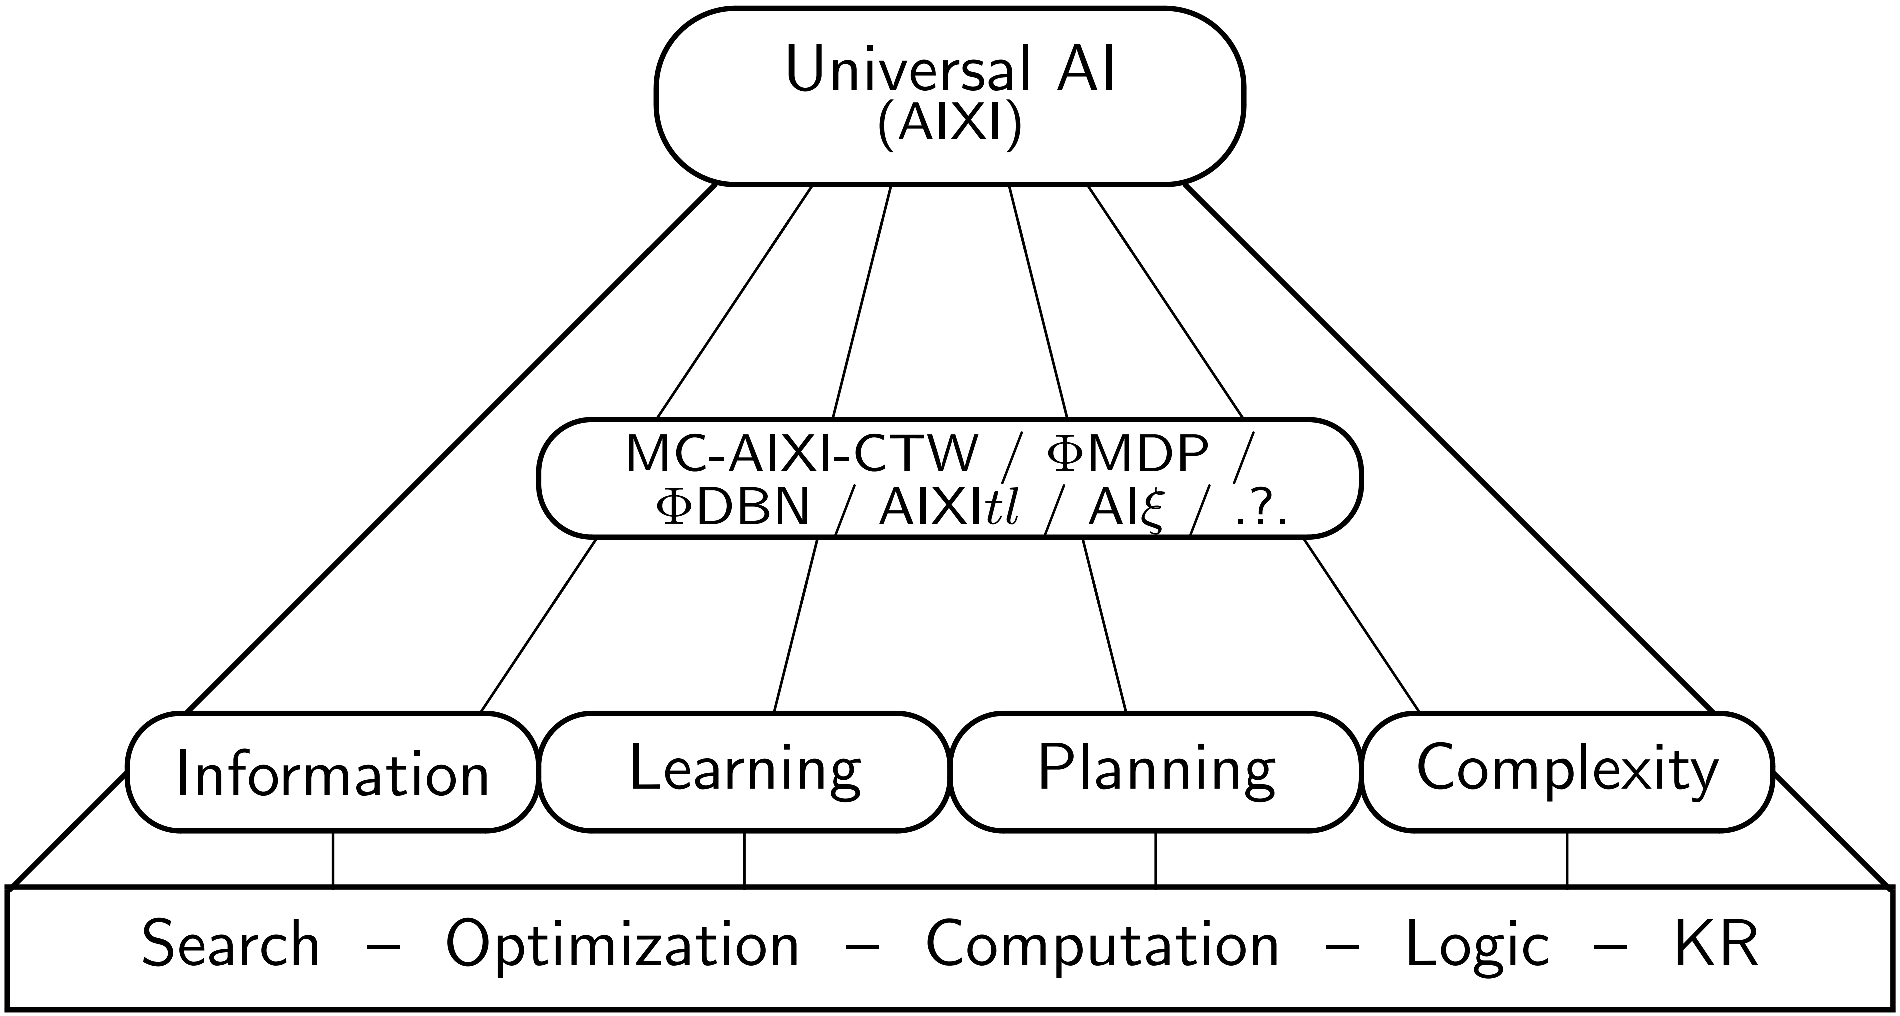

Agents $=$ General Framework, Interface $=$ Robots,Vision,Language

Temporal Difference without a Learning Rate

- Reinforcement learning TD update:

$$V_s^{t+1} = V_s^t + E_s^t \beta_t(s, s_{t+1})(r_t + \gamma V_{s_{t+1}}^t - V_{s_t}^t) \ \forall s$$

- Learning rate β_t and eligibility trace E_s^t from stat. principles.

- Normally: Learning rate α, free parameter, tuned by hand.

- $\beta_t(s, s_{t+1}) \ := \ \frac{1}{N_s^t} \cdot \frac{1}{1 - \gamma E_{s_{t+1}}^t / N_{s_{t-}}^t}$

- In every setting that we have tested,

superior performance & fewer parameters to tune

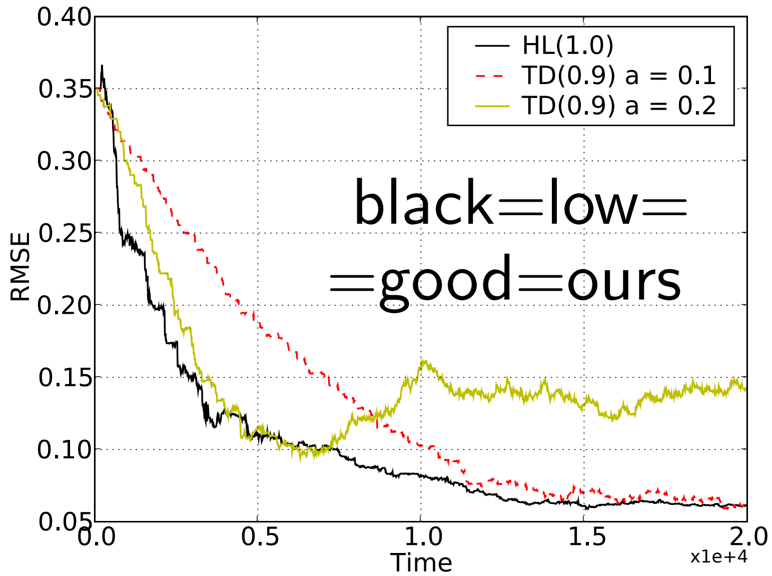

[HL'07]

Predictive Hypothesis Identification

- Given: $D\equiv(x_1,...,x_n)$ sampled from $p(D|\theta_0)$ with unknown $\theta_0\in\Omega$.
- Goal: Predict m future obs. $\boldsymbol{x}\equiv(x_{n+1},...,x_{n+m})$ well.
- Approach: Summarize data D by composite hypothesis $H\subseteq\Omega$ (point | finite set | interval | convex set)
- If θ_0 is true parameter, then $p(\boldsymbol{x}|\theta_0)$ is obviously the best prediction.
- If θ_0 unknown, then the Bayesian predictive distribution $p(\boldsymbol{x}|D)=\int p(\boldsymbol{x}|\theta)p(\theta|D)d\theta=p(D,\boldsymbol{x})/p(D)$ is best.
- Approx. full Bayes by predicting with hypothesis H, i.e.
- Use (comp) likelihood $p(\boldsymbol{x}|H)=\frac{1}{P[H]}\int_H p(\boldsymbol{x}|\theta)p(\theta)d\theta$ for prediction.
- The closer $p(\boldsymbol{x}|H)$ to $p(\boldsymbol{x}|\theta_0)$ or $p(\boldsymbol{x}|D)$ the better H's prediction.
- Measure closeness with some distance function $d(\cdot,\cdot)$.
- Since $\boldsymbol{x}$ and θ_0 are unknown, we must sum or average over them.
- $H^{\text{PHI}}:=\arg\min_H\int\int d(p(\boldsymbol{x}|H),p(\boldsymbol{x}|\theta))\,p(\theta|D)d\boldsymbol{x}\,d\theta$
- Satisfies many desirable properties (inv.,≈ML,≈MAP,eff.). [H'08]

The Loss Rank Principle for Model Selection

Let $\hat{f}_D^c : \mathcal{X} \to \mathcal{Y}$ be the (best) regressor of complexity c on data D.

The loss *Rank* of $\hat{f}_D^c$ is defined as the number of other (fictitious) data D' that are fitted better by $\hat{f}_{D'}^c$ than D is fitted by $\hat{f}_D^c$.

- c is small $\Rightarrow$ $\hat{f}_D^c$ fits D badly $\Rightarrow$ many other D' can be fitted better $\Rightarrow$ *Rank* is large.
- c is large $\Rightarrow$ many D' can be fitted well $\Rightarrow$ Rank is large.
- c is appropriate $\Rightarrow$ $\hat{f}_D^c$ fits D well *and* not too many other D' can be fitted well $\Rightarrow$ *Rank* is small.

LoRP: Select model complexity c that has minimal loss *Rank* [H'07]

Unlike most penalized maximum likelihood variants (AIC,BIC,MDL),

- LoRP only depends on the regression and the loss function.
- It works without a stochastic noise model, and
- is directly applicable to any non-parametric regressor, like kNN

Distribution of Mutual Information (MI)

- θ_{ij} = joint prob. of $(i,j) \in \{1,...,r\} \times \{1,...,s\}$. Marginals θ_{i+}, θ_{+j}
- Widely used measure for (in)dependence of random vars $\imath$ and $\jmath$ is: Mutual Information: $I(\boldsymbol{\theta}) \;=\; \sum_{i=1}^{r} \sum_{j=1}^{s} \theta_{ij} \log \frac{\theta_{ij}}{\theta_{i+}\theta_{+j}}$
- Applications are abundant, e.g. connecting nodes in Bayesian Nets.
- Problem: θ_{ij} unknown $\Rightarrow$ freq. est. from data: $\theta_{ij} \approx \hat{\theta}_{ij} := \frac{n_{ij}}{n}$
- Problems of Point Estimate: $I(\hat{\boldsymbol{\theta}})$ gives no information about its accuracy, e.g. $I(\hat{\boldsymbol{\theta}}) \neq 0$ true dependency -or- random fluctuation.
- Bayesian Solution: Compute posterior of MI based on (Dirichlet) prior over θ_{ij}: $p(I|\boldsymbol{n}) = \int \delta(I(\boldsymbol{\theta}) - I) p(\boldsymbol{\theta}|\boldsymbol{n}) d^{rs}\boldsymbol{\theta}$
- Compute by systematic expansion in $1/n$ (fast & accurate) [H'01]
- Extension to Missing Data (*global* max by EM) [HZ'03-05]

Exact Bayesian Inference on Infinite Trees

- Given: i.i.d. data from an unknown distributio

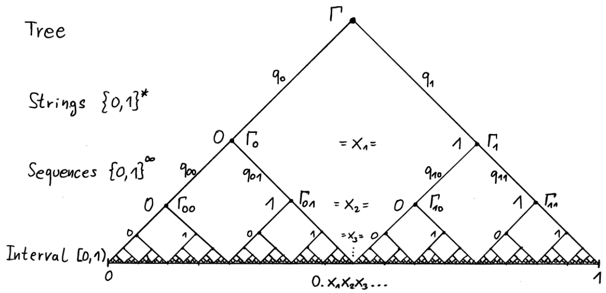

- Goal: predict future items or distribution.
- Partition domain: recursively ⇒ infinite tree
- BayesTree: Assign prior to "subdivide" ⇒ prior over ∞-trees [H'05]
- Algorithm: exact, fast, simple! for posterior, data evidence, predictive distribution, effective model dimension, ... [H'05]

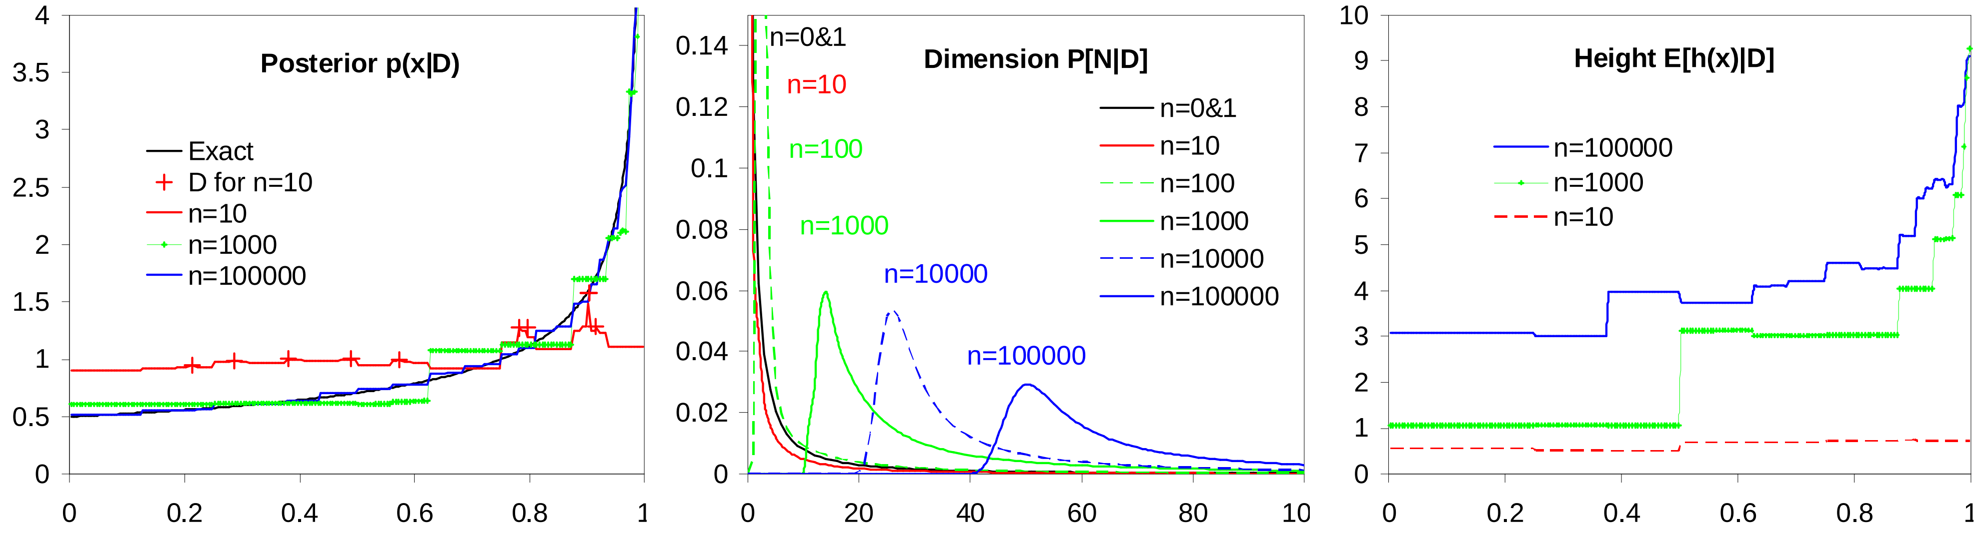

Robust Bayes and Imprecise Probabilities

- Problem: Bayesian prior may not be known/knowable precisely.
- Solution: Choose set of (all) reasonable priors $\Rightarrow$ robustness
- Set of priors $\Rightarrow$ set of posteriors $\Rightarrow$ set of optimal decisions.
- Imprecise Dirichlet Model=Set of Dirichlets $= \{\text{Diri}(\alpha) : \sum \alpha_i \leq s\}$
- IDM satisfies symmetry principle and is reparametrization invariant.
- Derivation of exact, conservative, and approximate, robust and credible interval estimates for a large class of statistical estimators, including the entropy and mutual information. [H'03]
- Further: general error propagation formulas, IDM for product spaces, robust credible sets. [H'03]
- Application: Robust inference of dependency-trees. [ZH'03]

Bayesian Change Point Detection

- Given: Very noisy data sequence $\boldsymbol{y} = (y_1, ..., y_n)$ with jumps.
- Goal: Find (regress) underlying piecewise constant function.
- Estimate segment levels $\boldsymbol{\mu} = (\mu_1, ..., \mu_k)$, boundaries $\boldsymbol{t} = (t_0, ..., t_k)$, their number k, and hyper-parameters.
- Bayesian regression: Compute posterior $P(\boldsymbol{\mu}, \boldsymbol{t}, k|\boldsymbol{y})$ and mean/MAP summaries from prior and likelihood.
- Algorithm based on DP: exact, fast, simple. [H'07,Lindley Prize]

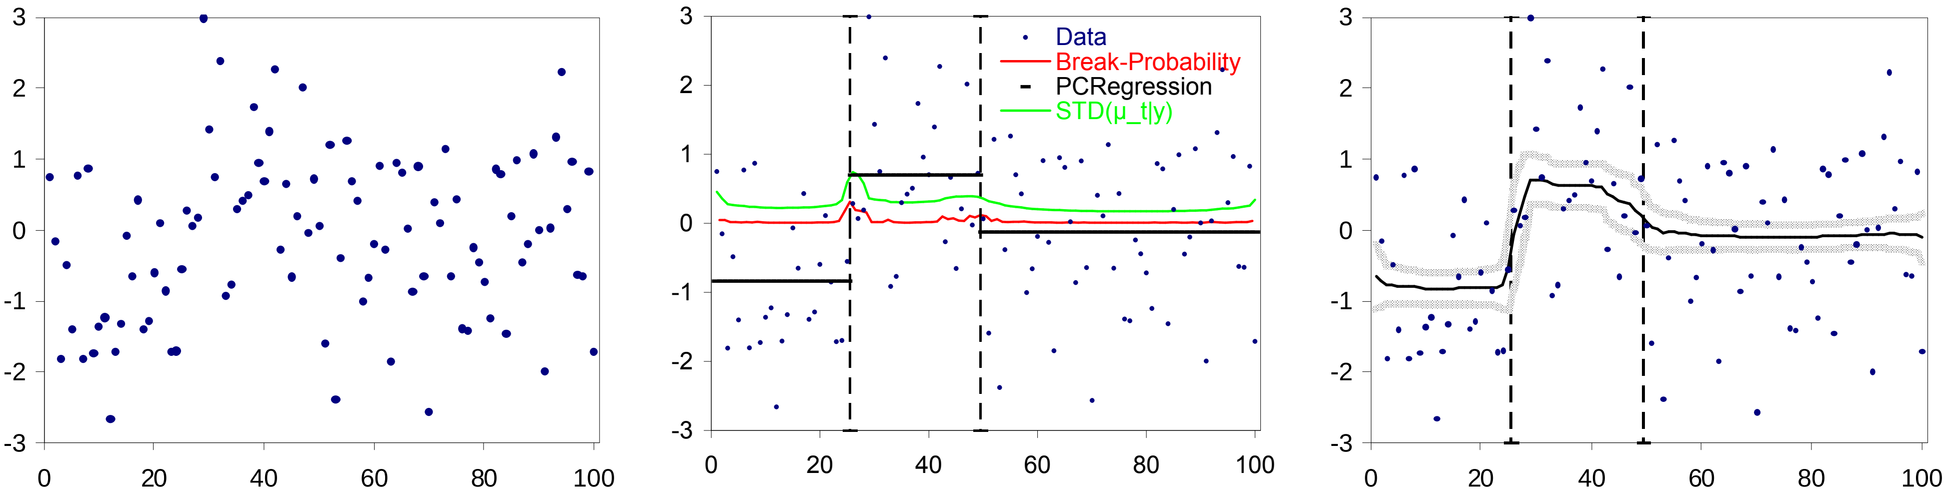

DNA Copy Number Estimation

- The Copy Number (CN) of a genomic region is the number of the copies of DNA in that region (CN=2 in a healthy cell).
- Tumors affect the DNA CN in the genome.
- CN can be modeled as piecewise constant function along the genome.
- A very noisy signal of CN can be measured by DNA MicroArrays.
- Exact Bayesian Piecewise Constant regression (BPCR) can be used to reconstruct the true CN. [H'07,Lindley Prize,RHBK'08-10]
- Helps in finding cancer-relevant genes.

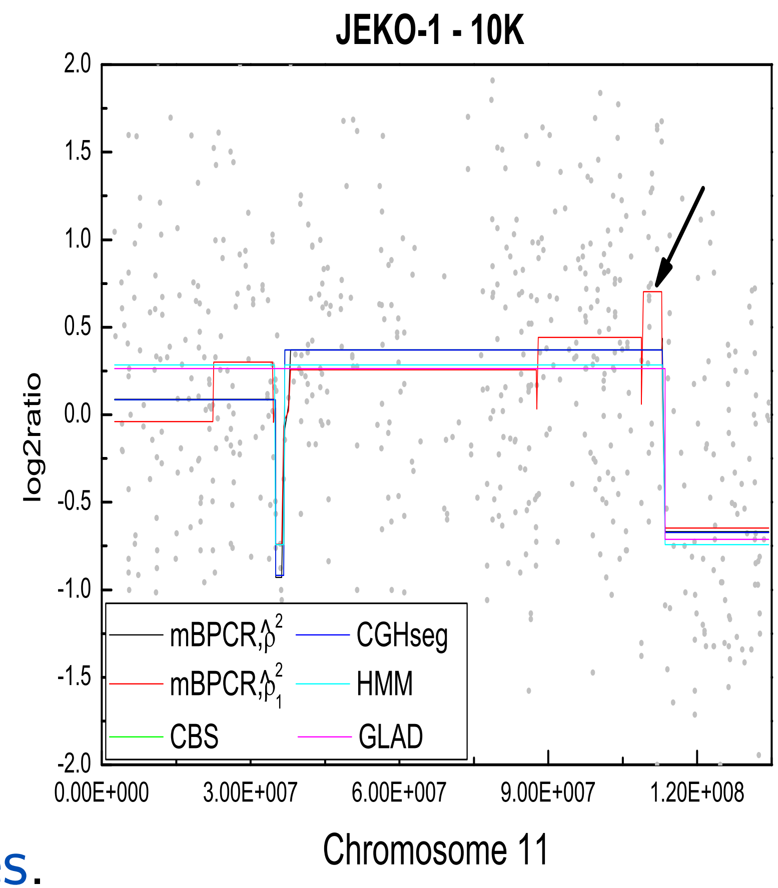

Bayesian Sequence Prediction

- Predict with known (subjective) Bayes mixture $\xi(x_1...x_n) := \sum_{\nu\in\mathcal{M}} w_\nu \nu(x_1...x_n)$ in place of unknown (objective) true distribution μ.

- Bound on the relative entropy $\text{KL}(\mu||\xi) \leq \ln w_\mu^{-1} < \infty$ [S'78]

$\Rightarrow \xi(x_t|x_1...x_{t-1}) \to \mu(x_t|x_1...x_{t-1})$ rapid conv. with prob. 1.

- ξ is also optimal in a decision-theoretic sense w.r.t. any bounded loss function: $\text{Loss}_{1:n}(\xi)/\text{Loss}_{1:n}(\mu) \to 1$ [H'01,H'03]

- No structural assumptions on model class $\mathcal{M}$ and $\nu \in \mathcal{M}$!

- More results: Fast convergence, optimality, continuous $\mathcal{M}$, multi-step predictions, similar but weaker for MDL, ... [H'04]

- Main assumption: Sequence $x_1 x_2 x_3...$ sampled from $\mu \in \mathcal{M}$.

Philosophical Questions

- How to choose the prior (w_ν)?
- How to choose the model class $(\mathcal{M})$?
- When is an individual sequence random?
- What does probability mean?

[S'64,ML'66,H'04,H'06]

Minimum Description Length Principle

- Probability axioms give no guidance of how to choose the prior.
- Occam's razor is the only general (always applicable) principle for determining priors, especially in complex domains typical for AI.
- Prior $= 2^{-\text{descr.length}}$ — Universal prior $= 2^{-\text{Kolmogorov complexity}}$.
- Prediction $\widehat{=}$ finding regularities $\widehat{=}$ compression $\widehat{=}$ MDL.
- MDL principle: from a model class, a model is chosen that: minimizes the joint description length of the model and the data observed so far given the model.
- Similar to (Bayesian) Maximum a Posteriori (MAP) principle.
- MDL often as good as Bayes but not always. [PH'04..PH'06..H'09]

Prediction with Expert Advice

PEA combines predictions of Experts $i \in \{1,..,n\}$: Two major variants:

- WM: $\text{Prob}[I_t^{\text{WM}} = i] \propto \exp[-\eta_t \cdot \text{Loss}_{<t}(\text{Expert}_i) - k^i]$
- FPL: $I_t^{\text{FPL}} = \arg\min_i \{\text{Loss}_{<t}(\text{Expert}_i) + (k^i - \text{Random}_t^i)/\eta_t\}$

Notation: $x_{<t} := (x_1, ..., x_{t-1})$ and $\boldsymbol{y}_t = (y_t^1, ..., y_t^n)$.

For $t = 1, 2, ..., T$

- Predict $y_t^{\text{PEA}} := \text{PEA}(x_{<t}, \boldsymbol{y}_t, \text{Loss})$
- Observe $x_t := \text{Env}(\boldsymbol{y}_{<t}, x_{<t}, y_{<t}^{\text{PEA}})$
- Receive $\text{Loss}_t(\text{Expert}_i) := \text{Loss}(x_t, y_t^i)$ for each Expert ($i = 1, ..., n$)
- Suffer $\text{Loss}_t(\text{PEA}) := \text{Loss}_t(x_t, y_t^{\text{PEA}})$

No statistical assumption on sequence!, any bounded loss function, (in)finite number of experts, adaptive learning rate η_t.

Results: $\text{Loss}_{1:T}(\text{PEA})/\text{Loss}_{1:T}(\text{Expert}_i) \to 1 \quad \forall i$ [LW'89..H'04]

Algorithmic Probability & Information Theory

Kolmogorov complexity $K(x) := \min\{\ell(p) : U(p) = x\}$ is a/the universal (domain independent) measure of the information content of x.

Properties: Optimal compressor, finds all effective regularities.

Efficient approximations: Shannon entropy, specific MDL codings, Lempel-Ziv compression, ...

Applications: Universal similarity metric (phylogeny, language, music, astronomy) [CV'05]

Solomonoff's universal a priori probability

$M(x) := \sum_{p:U(p)=x*} 2^{-\ell(x)} \approx 2^{-K(x)}$ assigns high/low probability to simple/complex strings, thus quantifying Occam's razor.

Applications: Optimal and universal sequential predictions. [S'64..H'04]

The Fastest Algorithm for All Problems

- Let $p^*: X \to Y$ be a given algorithm or problem specification.
- Let p be any algorithm, computing provably the same function as p^* with computation time provably bounded by the function $t_p(x)$.
- Then the algorithm M_{p^*} computes $p^*(x)$ in time

$$\boxed{time_{M_{p^*}}(x) \;\leq\; 5\cdot t_p(x) + \text{lower-order-terms}}$$

- Neither p, t_p, nor the proofs need to be known in advance for the construction of $M_{p^*}(x)$

[H'01]

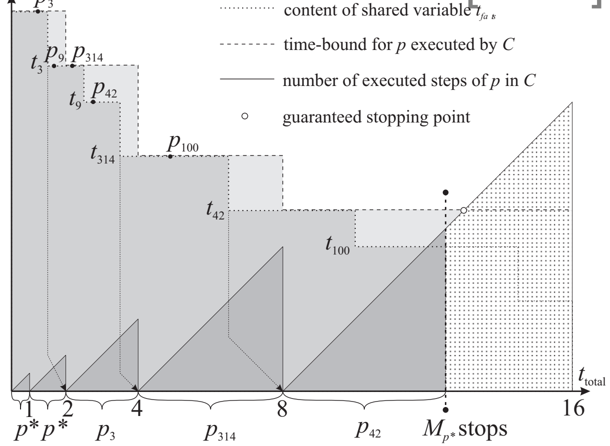

- Idea: Enumerate all p provably equivalent to p^* and execute the currently fastest one.
- Catch: Lower order terms are huge and dominate in practice.

Optimization

- Linear time approximation algorithm for Knapsack problem. [MH'02]
- New fitness uniform selection scheme for GA. [H'02..HL'06]
- Exploration versus exploitation (Opt,Bayes,PAC,Asymp) [H'00,P'04,RH08,LH'11,...]

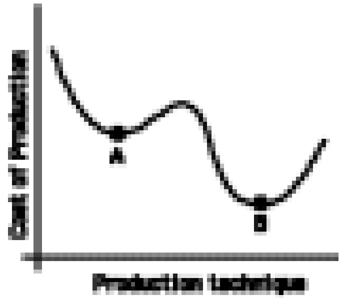

Image Processing and Computer Vision

- Image-Based Car Damage Detection (ICAR) [2007-2013]
- 2D-3D Model-Based Image Registration [HB'09,JYH'10,JHB'11]

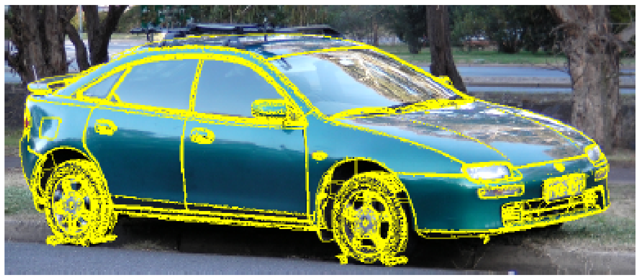

Computer Vision and Graphics (1996–2000)

- Stereo-Vision for uncalibrated and non-linear infrared cameras to sub-millimeter precision.

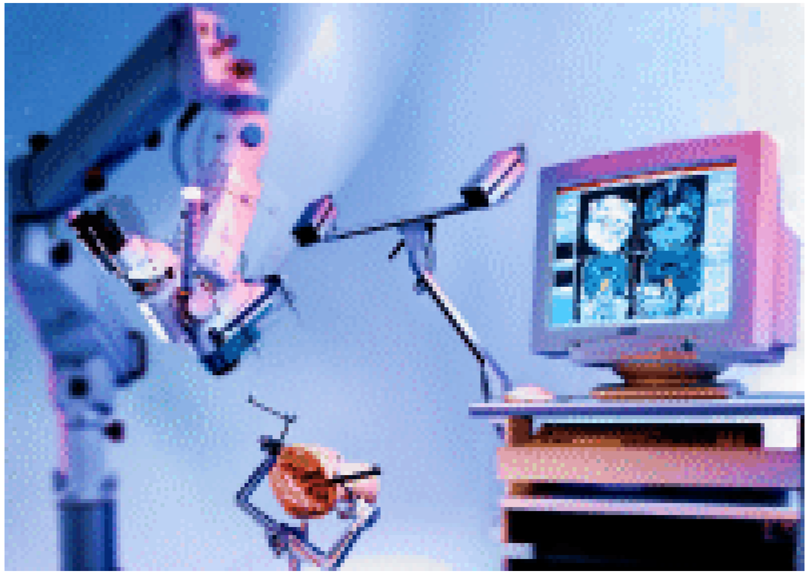

- Real-time software volume renderer.
- Automatic CT/MRI/Ultrasound image fusion and 3D surface matching.
- Planning and steering of robotic microscope tripods.
- Image and volume segmentation (e.g. brain/bone/tissue/water).
- Image and volume data enhancement and post-antialiasing algorithms based on finite-element interpolation (patented).

[1996-2000, mostly confidential/unpublished]

More Industrial R&D (1996–2000)

- Brachytherapy planning system.
- A dose algorithm (PencilBeam) for radiotherapy for IMRT.

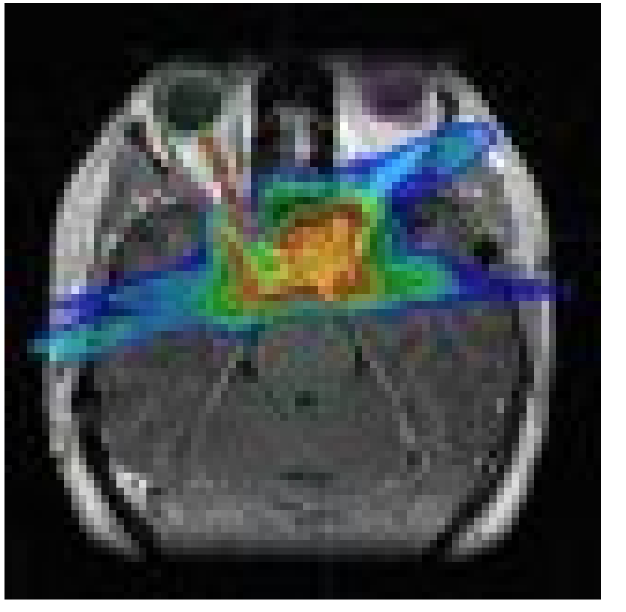

[1996-2000, confidential/unpublished]

Particle Physics (1993-1996)

- Motivation: Theory of Everything
  $\stackrel{?}{=}$ String Theory = Gravity + Quantum Theory.
- Will it also solve the problem of the interpretation of quantum theory? (Schrödinger's cat)
- Addressing the proton spin problem (unsuccessful) [H'95]
- Computation of the gluon, quark, and meson correlators and masses in the instanton liquid model. [H'93-97]
- PhD in non-perturbative QuantumChromoDynamics, in particular instantons in QCD. [H'96]
- Explanation of the exponential fermion mass spectrum between successive generations. [BH'97]

[1993-1996, PhD]

Early Student Work (1983-1992)

- Implementation of a classifier system, allowing for comparison of many popular variants. [H'92]
- Proof of equivalence of ranking and tournament selection. [H'92]
- Reinforcement feedback for unsupervised learning Hebb nets. [H'90]
- THE CAD program for 8 bit computers (in Assembler!). [H'87]
- Miscellaneous: Implementation of
  - a member organization program in DBase, [H'83]
  - a user interface for an expert system under GEM, [H'87]
  - a protection module and organizer for licensing programs. [H'93]

Generally Accessible Papers

about my foundational U(A)I work (2000-2010)

[Hut12] M. Hutter. One decade of universal artificial intelligence. In *Theoretical Foundations of Artificial General Intelligence*. Atlantis Press, 2012.

[RH11] S. Rathmanner and M. Hutter. A philosophical treatise of universal induction. *Entropy*, 13(6):1076–1136, 2011.

[Hut11] M. Hutter. Algorithmic randomness as foundation of inductive reasoning and artificial intelligence. In *Randomness through Computation*, chapter 12, pages 159–169. World Scientific, 2011.

[Hut10] M. Hutter. A complete theory of everything (will be subjective). *Algorithms*, 3(4):329–350, 2010.

[LH07] S. Legg and M. Hutter. Universal intelligence: A definition of machine intelligence. *Minds & Machines*, 17(4):391–444, 2007.

[Hut07] M. Hutter. Algorithmic information theory: a brief non-technical guide to the field. *Scholarpedia*, 2(3):2519, 2007.

Thanks! Questions? Details:

Current Group Members (including PhD students):
Peter Sunehag - Wen Shao - Mayank Daswani - Di Yang - Tor Lattimore - Phuong Nguyen - Srimal Jayawardena - Matthew Robards.

Thanks to all of them and also
to all collaborators and past group members.

Financial Support (Grants) SNF, ARC, CE.

Book on Universal Artificial Intelligence $\Longrightarrow$

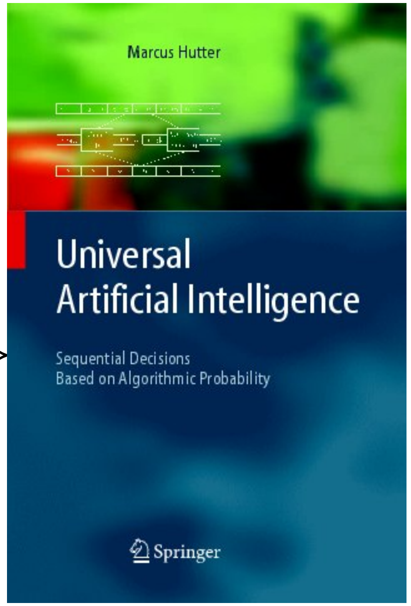

Detailed references: www.hutter1.net/official/publ.htm

Projects at www.hutter1.net/ai/projects.htm

Compression competition with 50'000 Euro prize at prize.hutter1.net

Jobs: For PostDoc and PhD positions at RSCS and NICTA, Australia, see www.hutter1.net/official/jobs.htm