Cooperative man-machine system modeling

Alfred Hero

University of Michigan - Ann Arbor

June 17, 2013

Outline

Humans and Machines

Mumford-Wiener view of evolving human-machine relationship

- pretechnic (pre-1000): manual labor, little or no mechanization.
- eotechnic (1000): machine assists man to do labor.
- paleotechnic (1700): man assists machine to do labor.
- neotechnic (1900): semi-automated machine does labor.
- cybertechnic (1940): fully automated self-regulated machines

Lewis Mumford 1934, *Technics and Civilization*
Norbert Wiener 1948, *Cybernetics*

Pretechnic wheat harvesting: winnowing

fork winnowing

wind winnowing

sources: `osirisnet.net`. `www.4to40.com`

Eotechnic/Paleotechnic: mechanically assisted harvesting

McCormick reaper

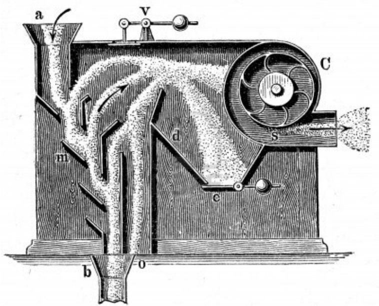

winnowing machine

sources: `www.fasttrackteaching.com/burns/`. Encyclopedia of education, St. Peterburg, 1896.

Paleotechnic/Neotechnic: towards full automation

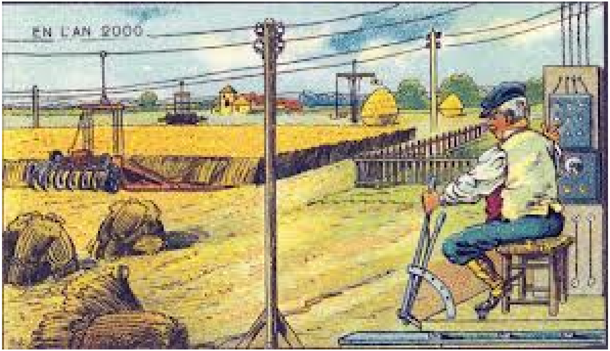

Villemard's vision (1910)

commons.wikimedia.org

Progress on agricultural harvesting automation

Combine harvestor (1930's) Autonomous harvestor (2011)

`commons.wikimedia.org, www.popsci.com`

Progress on agricultural harvesting automation

Combine harvestor (1930's) Autonomous harvestor (2011)

`commons.wikimedia.org, www.popsci.com`

Computer-assisted information retrieval systems - a similar path of progress?

Computer-assisted information retrieval systems

Fit into Mumford's four stages of development:

- pre-technic (pre-1500AD): Small body of accessible knowledge. Few persons have access to information.
- eotechnic (1500AD): Larger bodies of accessible knowledge. Universal encyclopedias of knowledge created in 18th century. Public library collections.
- paleotechnic (1970AD): Great bodies of accessible knowledge. On-site electronic card catalogs.
- neotechnic (1990AD): Massive bodies of accessible information. Online indexing and retrieval.

Mumford's model: progressive supplantation of human labor:

man $\Rightarrow$ machine-aided man $\Rightarrow$ man-aided machine $\Rightarrow$ automaton

Emerging trend: re-engagement of human in mechanism

Examples of human re-engagement in information retrieval systems

- Reputation-based webpage ranking: Google's pagerank algorithm, Kleinberg's HITS algorithms

Emerging trend: re-engagement of human in mechanism

Examples of human re-engagement in information retrieval systems

- Reputation-based webpage ranking: Google's pagerank algorithm, Kleinberg's HITS algorithms
- Socio-personalized recommendation systems: use of social, preferential, geographical, and behavioral user-profiles

Emerging trend: re-engagement of human in mechanism

Examples of human re-engagement in information retrieval systems

- Reputation-based webpage ranking: Google's pagerank algorithm, Kleinberg's HITS algorithms
- Socio-personalized recommendation systems: use of social, preferential, geographical, and behavioral user-profiles
- Electronic fraud detection: socio-temporal inter-connectivity analysis

Emerging trend: re-engagement of human in mechanism

Examples of human re-engagement in information retrieval systems

- Reputation-based webpage ranking: Google's pagerank algorithm, Kleinberg's HITS algorithms
- Socio-personalized recommendation systems: use of social, preferential, geographical, and behavioral user-profiles
- Electronic fraud detection: socio-temporal inter-connectivity analysis
- Crowdsourced semantic tagging/annotation: hire many people to hand annotate (Amazon Mechanical Turk, ImageNet).

Emerging trend: re-engagement of human in mechanism

Examples of human re-engagement in information retrieval systems

- Reputation-based webpage ranking: Google's pagerank algorithm, Kleinberg's HITS algorithms
- Socio-personalized recommendation systems: use of social, preferential, geographical, and behavioral user-profiles
- Electronic fraud detection: socio-temporal inter-connectivity analysis
- Crowdsourced semantic tagging/annotation: hire many people to hand annotate (Amazon Mechanical Turk, ImageNet).
- Interactive interfaces: speech, haptic, neural, visual, virtual environments (Google glass)

This talk pursues two specific directions:

- Collaborative human-in-the-loop systems: interactive real-time cooperative search, localization, detection, classification tasks.
- Multidimensional query representations: partial orders, Pareto selection

Outline

Computational models for information extraction

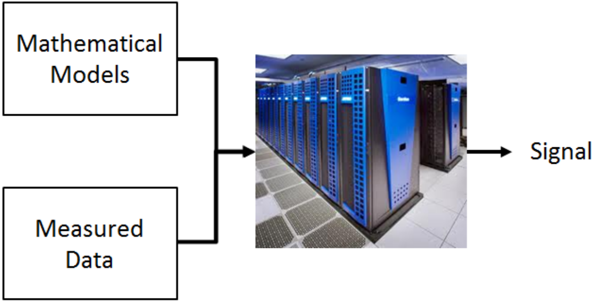

A human designs the function f that computer evaluates on data and models:

$$\text{signal} = f(\text{data}, \text{models})$$

Old model: human operates machine

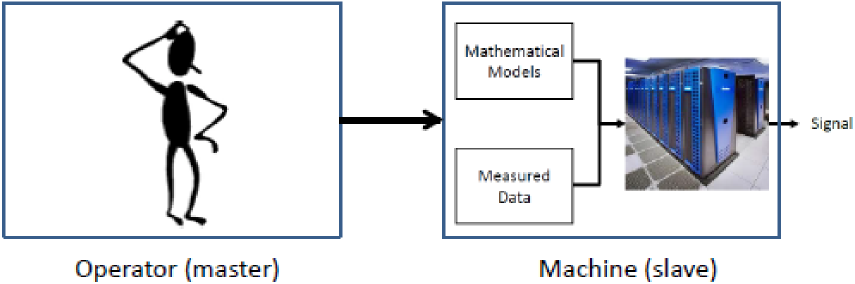

Relationship

Human is master - Machine is slave

Computational models for information extraction

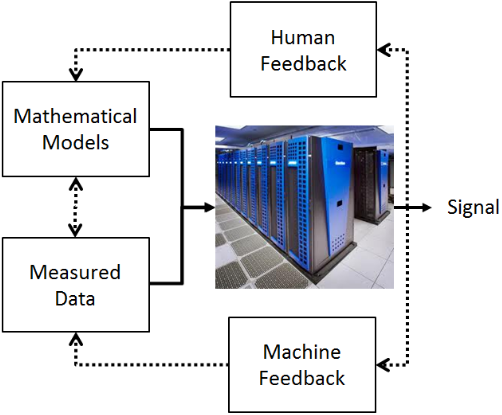

Feedback refines and improves model and data acquisition

- Human-assisted processing, relevance feedback learning
- Plan-ahead sensing, sensor management, sequential DOE

New model: human operates cooperative human-machine

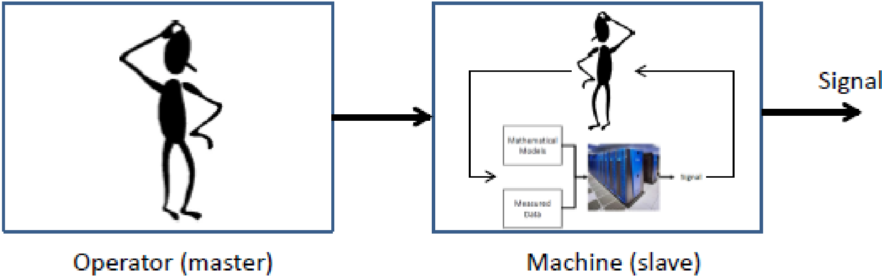

Relationship

Human is master - Computer with human-in-the-loop is slave

Two questions of importance

- Fundamental design principles for collaborative human-machine systems
  - Impedance matching: load balancing, compatible I/O?
  - Accounting for human latency, delay, accuracy, responsiveness?
  - System testing and reliability?

Two questions of importance

- Fundamental design principles for collaborative human-machine systems
  - Impedance matching: load balancing, compatible I/O?
  - Accounting for human latency, delay, accuracy, responsiveness?
  - System testing and reliability?
- Relevant mathematical models
  - Must account for dynamics of human-machine interaction?
  - Must be robustly predictive of performance?
  - Must be task specific?

Two questions of importance

- Fundamental design principles for collaborative human-machine systems
  - Impedance matching: load balancing, compatible I/O?
  - Accounting for human latency, delay, accuracy, responsiveness?
  - System testing and reliability?
- Relevant mathematical models
  - Must account for dynamics of human-machine interaction?
  - Must be robustly predictive of performance?
  - Must be task specific?

Successful systems will likely exploit complementary strengths of human and machine.

Human factors and man-machine interfaces

Function allocation: assign task to computer OR human

Human-machine cooperation: assign it to computer AND human

Dimensions of human performance (Laughner et al 2006)

- Attention: ability to attend to a stimulus for extended period
- Perception: ability to detect specific stimulus patterns
- Psychomotor skill: ability to maintain system stability
- Physical skill: ability to accomplish sustained, effortfull work
- Cognitive skill: ability to apply concepts and rules to information

Comprehensive human model is probably futile

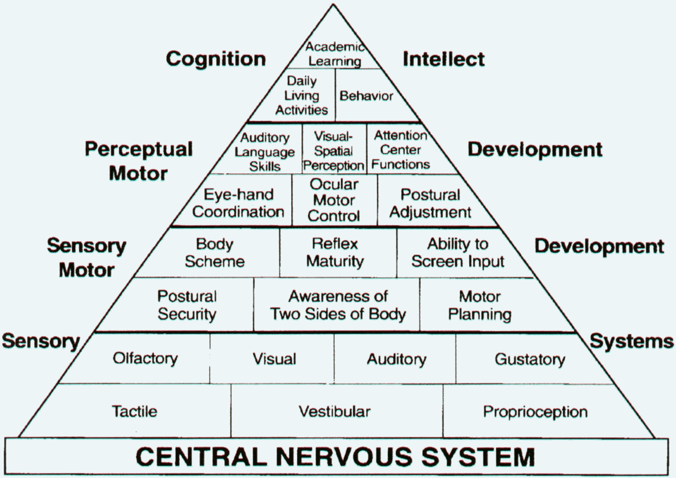

Figure 5. Pyramid of Learning. (Williams & Shellenberger, 1-4)

Human-computer complementarity

Computer information processing (CIP):

- Algorithmic determinism: mathematical and reproducible
- Fast: many operations in short time period.
- Reliable: does not tire or get stressed out.
- Inexpensive: dollars per flop
- Numerically accurate: gives results to n-th decimal place

Human-computer complementarity

Computer information processing (CIP):

- Algorithmic determinism: mathematical and reproducible
- Fast: many operations in short time period.
- Reliable: does not tire or get stressed out.
- Inexpensive: dollars per flop
- Numerically accurate: gives results to n-th decimal place

Human information processing (HIP):

- Heuristic intuition: non-linear reasoning with flashes of insight
- Contextual: can filter out irrelevant information
- Adaptable: mixes deductive, inductive and heuristic reasoning.
- Concept generalization: extrapolate from current evidence, make educated guess at a solution and work backwards

Human factors and degradation functions

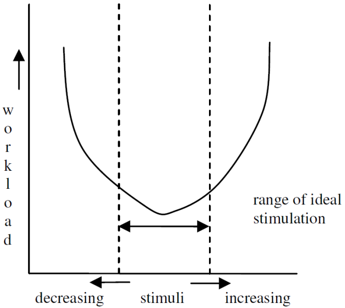

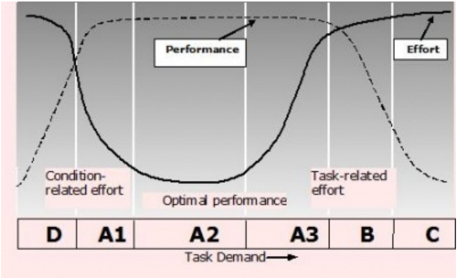

Under vs over stimulation

Effort and performance

Degradation function relative to a stressor:

- Stimulation level and workload (effort)
- Time on task (fatigue)
- Attention (distractions in environment)

A few relevant mathematical models

- Mathematical psychology: (Wagenmakers 2007)
  - Simplified Ratcliff diffusion model for human response accuracy and speed ($y = a\nu$)

$$P_c = \frac{1}{1+\exp(-y)}, \quad \mathrm{MDT} = \left(\frac{a}{2\nu}\right)\frac{1-\exp(-y)}{1+\exp(-y)}$$

  - Bayesian models for human decisionmaking on bandit problems (Steyvers 2009, Gibson 2011).
- Educational testing: (Wang 2013)
  - Rasch's dynamic item response model:

$$P(X_{il} = 1|\theta_i, d_l) = \frac{\exp(\theta_i - d_l)}{1+\exp(\theta_i - d_l)}$$

- Mathematical economics (von Neuman and Morgenstern, 1944)
  - Behavioral game theory and team theory models (Kim 1987).

Ratcliff's diffusion model for evidence accumulation

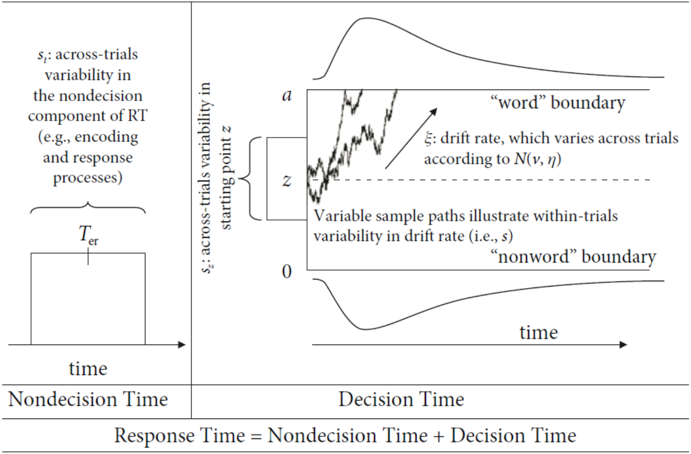

Figure 1. Diffusion model account of evidence accumulation in the lexical decision task (see Ratcliff et al., 2004).

Outline

Example: collaborative target localization

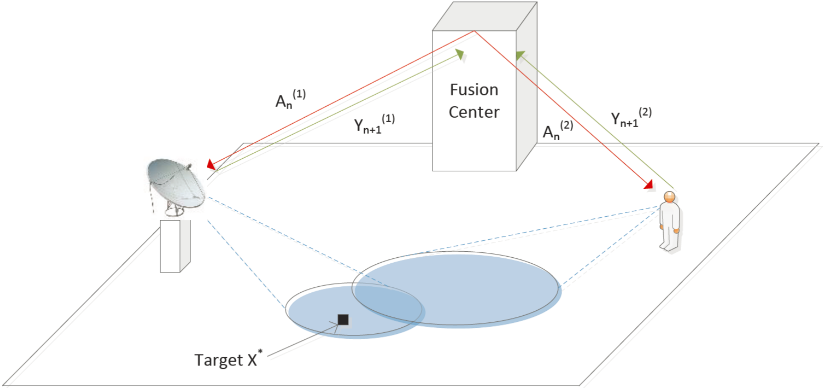

Tsiligkaridis and H, "A collaborative 20 questions model for target search with human-machine interaction," arXiv:1306.1922, 2013.

Markov decision process (MDP) framework

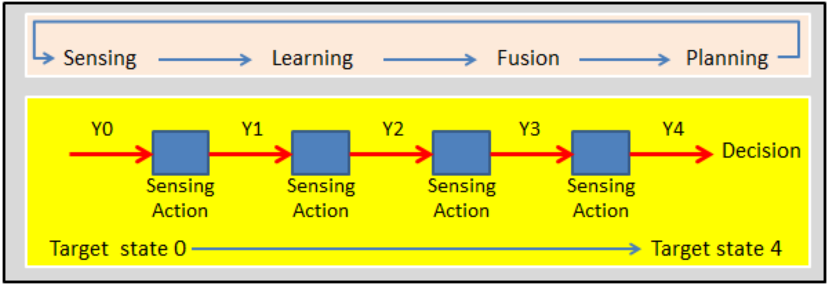

Multistage Markov decision process (MDP) (Bellman 1957)

- Human assistance can occur at any stage
- Full multistage optimization of MDP is intractible
- Useful framework for obtaining bounds and inspiration

Collaborative target localization model

Tsiligkaridis' (2013) proposed approach combines:

- Multi-player extension of "20 questions with noise" (Jedynak 2012)
  - MDP formulation with entropic reward function
  - Binary response with a BSC reliability function
  - Optimality of greedy probabilistic bisection policy (Horstein 1963)
- Spatial-resolution-limited human noise model (Jamieson 2012)
  - Human's context-integration advantage at coarse scales
  - Human's accuity disadvantage at fine scales
  - This noise model model is used as reliability function for each BSC

Related approaches

Our formulation is related to

- Interval discrete probabilistic bisection (Castro 2007, Burnashev&Zigangirov 1974)
- Noisy generalized binary search (Nowak 2009)
- Stochastic approximation and root finding (Robbins&Munro 1951)
- Sequential design of experiments (Robbins 1952, Whittle 1988)

Joint collaborative target tracking

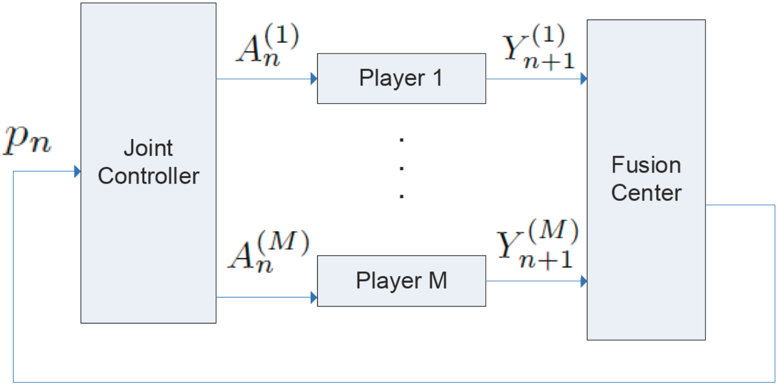

Tsiligkaridis and H, "A collaborative 20 questions model for target search with human-machine interaction," arXiv:1306.1922, 2013.

Cyclic coordinate-wise collaborative target tracking

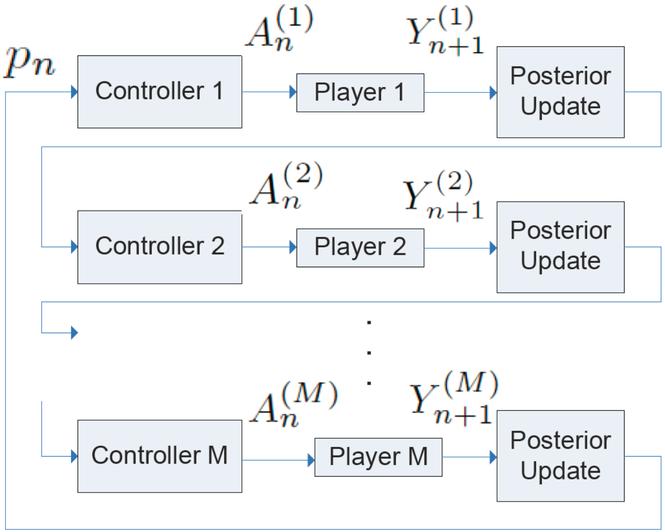

Tsiligkaridis and H, "A collaborative 20 questions model for target search with human-machine interaction," arXiv:1306.1922, 2013.

20 questions formulation of localization problem

Mathematical formulation

- Target state at unknown location $X^* \in \mathbb{R}^d$
- Controller query to agent m at time n:

$$\text{“}\texttt{Does } X^* \texttt{ lie in } A_n^{(m)} \subset \mathbb{R}^d\text{?”}$$

- Correct answer:

$$Z = I(X^* \in A_n^{(m)}) \in \{0, 1\}$$

- m-th agent's noisy response:

$$\text{“}Y_{n+1}^{(m)} = 0\text{” or “}Y_{n+1}^{(m)} = 1\text{”}$$

- Error probability (BSC):

$$P(Y_{n+1}^{(m)} \neq Z_n^{(m)} | Z_n^{(m)}) = \epsilon_m$$

- Posterior distribution of target location after N queries:

$$p_N(B) = P(X^* \in B | \mathcal{A}_N, \mathcal{F}_N), \qquad B \subset \mathbb{R}^d$$

Jedynak's theorem for case of $M = 1$ agent

Controller asks questions according to policy on filtration $\{\mathcal{F}_n\}_{n=1}^{20}$

$$\pi = (\pi_0, \pi_1, \ldots, \pi_{20})$$

over sets $A_i \subset \mathbb{R}^d$.

Controller's objective: policy should minimize entropy :

$$\min_{\pi} E^{\pi}[H(p_N)]$$

Jedynak's theorem for case of $M = 1$ agent

Controller asks questions according to policy on filtration $\{\mathcal{F}_n\}_{n=1}^{20}$

$$\pi = (\pi_0, \pi_1, \dots, \pi_{20})$$

over sets $A_i \subset \mathbb{R}^d$.

Controller's objective: policy should minimize entropy :

$$\min_{\pi} E^{\pi}[H(p_N)]$$

Theorem (Bisection - Jedynak 2012)

If there is only a single agent and $d = 1$, then the probabilistic bisection policy $P_n(A_n) = \int_{A_n} p_n(x)dx = 1/2$ is optimal and at each step n maximizes the mutual information:

$$I(Y_{n+1}, X^*|A_n, \mathcal{F}_n) = I(Y_{n+1}, X^*|A_n, p_n)$$

Extension for case of general $M \geq 1$ agents

Theorem (Separation - Tsiligkaridis 2013)

Let the M agents have conditionally independent responses $Y_n^{(m)}$ given X^*, past queries, and past responses. Then the cyclic coordinate-wise policy π^* is optimal and achieves entropy loss:

$$C = \sum_{m=1}^{M} C(\epsilon_m) = \sum_{m=1}^{M} (1 - h_b(\epsilon_m))$$

where $b(\epsilon) = -\epsilon \log \epsilon - (1-\epsilon)\log(1-\epsilon)$.
Furthermore, all optimal policies π^* satisfy

$$\int_B p_n(x)dx = 2^{-M}, \ \ \forall B \in \pi^*$$

.

Collaborative tracking rate bounds

Theorem (Error rate - Tsiligkaridis 2013)

Under the cooperative sequential query policy the conditional mean estimator X_n has MSE that converges with exponential rate:

$$a \exp\left(-n\overline{C}\right) \leq E[\|X^* - X\|_2^2] \leq b \exp\left(-\frac{2}{3}n\overline{C}\right)$$

where $\overline{C} = \sum_{m=1}^{M}(1/2 - \sqrt{\epsilon_m(1-\epsilon_m)})$.

Optimal collaborative policy interpretations

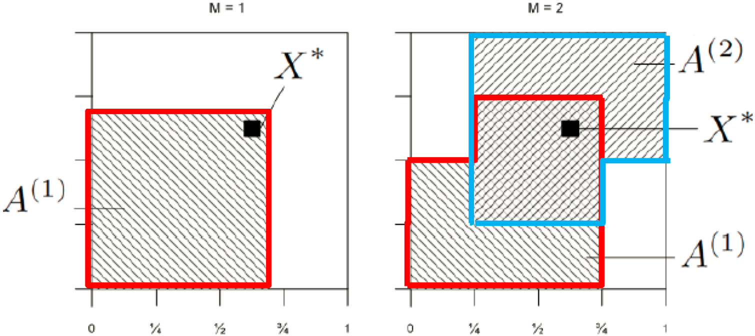

Two dimensional target optimal policies for 1 agent (left) and 2 agents (right)

- Jointly optimal policy reduces to simpler cyclic policy
- Joint policy does not decouple despite conditional independence of players
- Optimal player queries are overlapping but not identical

Human/machine error probability model (Jamieson 2012)

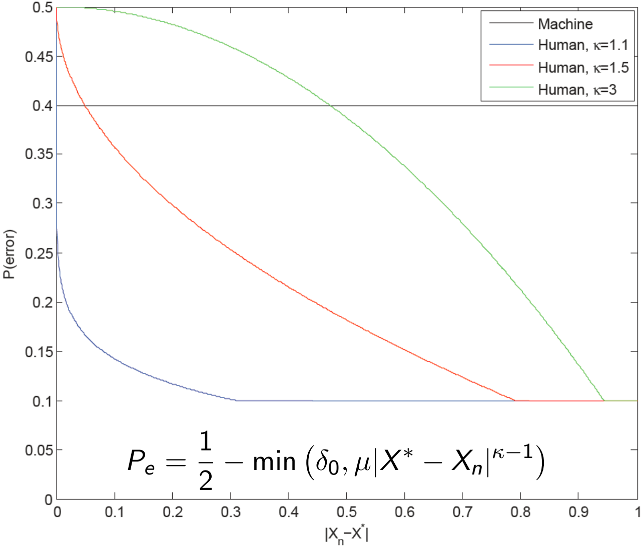

$$P_e = \frac{1}{2} - \min\left(\delta_0, \mu|X^* - X_n|^{\kappa-1}\right)$$

Human gain ratio for localization of 2D target

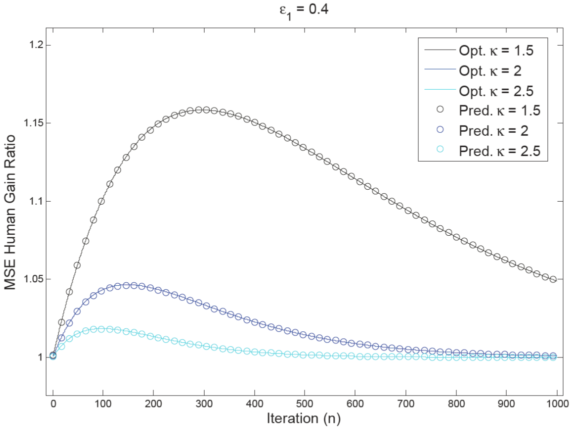

$$R_n(\kappa) \approx \frac{2^{-2/3} + 2^{1/3}}{2^{-2/3} + 2^{1/3}} \exp\left(-\frac{\mu^2}{30} (3 \cdot 2^{-7/3})^{2\kappa - 2} n e^{\overline{C}(\epsilon_1)} \frac{2n-1}{3} \right)$$

Open problems in 20 questions cooperative search

- Extension to case of unknown BSC crossover probabilities
- Inclusion of agent usage costs and/or agent switching costs
- Generalization to non-entropic objective functions: P_e, MSE
- Time critical applications: quickest detection, human response time models
- Other frameworks: decentralized collaborative networks, MABs, ARAP (Wei 2012)

Outline

1. Background
2. Cooperative human-in-the-loop systems
3. Example: collaborative sequential target localization
4. Multicriteria systems
5. Challenges

Single criteria search

Search a database, e.g. Google, for best matches to image query.

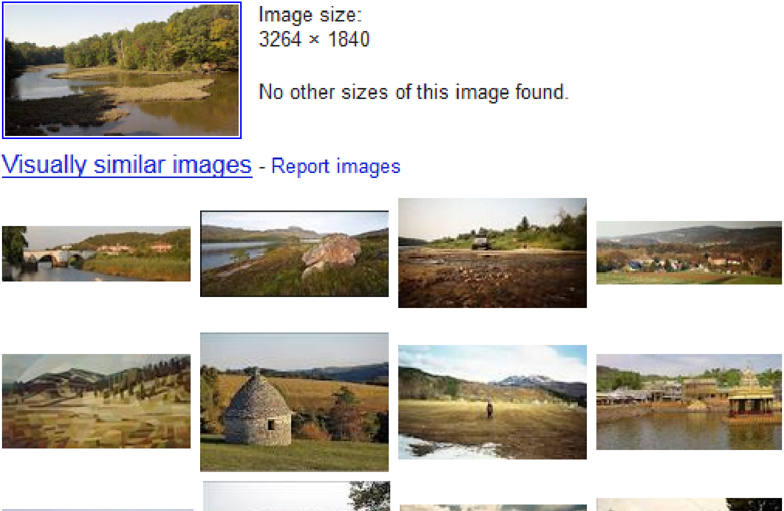

Single query search: linear ordering

Matches to query i sorted according to scalar dissimilarity

$$f(i_1) < f(i_2) < \ldots < f(i_n)$$

This encourages people to interact with Google's algorithm, leading to improvements

Single query search: linear ordering

Matches to query i sorted according to scalar dissimilarity

$$f(i_1) < f(i_2) < \ldots < f(i_n)$$

This encourages people to interact with Google's algorithm, leading to improvements

ML/SP/Imaging been dominated by algorithms that optimize *an* objective function

- Basis pursuit and dictionary learning find "a best match."
- Parametric estimation produces a ML, MAP, or min MSE estimator.
- Compressive sensing, matrix completion give "the best signal reconstruction."

Single query search: linear ordering

Matches to query i sorted according to scalar dissimilarity

$$f(i_1) < f(i_2) < \ldots < f(i_n)$$

This encourages people to interact with Google's algorithm, leading to improvements

ML/SP/Imaging been dominated by algorithms that optimize *an* objective function

- Basis pursuit and dictionary learning find "a best match."
- Parametric estimation produces a ML, MAP, or min MSE estimator.
- Compressive sensing, matrix completion give "the best signal reconstruction."

Emerging area in Machine Learning and SP: "Learning to rank"

Burges, Shaked, Renshaw, Lazier, Deeds, Hamilton, and Hullender, Learning to rank using gradient descent. In Proc. of the 22nd ICML, pp. 89-96. ACM, 2005.

Jamieson and Nowak, Active ranking using pairwise comparisons, arxiv, 2011.

Duchi, Mackey, and Jordan, The asymptotics of ranking algorithms, arxiv:1204.1688, Apr. 2012.

Dual query search

Sometimes a single query is not adequate
Alternative: dual query and multi-objective optimization

The similarity between query pair and database image i is now a vector $[f_1(i), f_2(i)]$.

One idea: rank according to scalarization $f_\lambda = \lambda f_1 + (1-\lambda) f_2$.

$$f_\lambda(i_1) < f_\lambda(i_2) < \ldots < f_\lambda(i_n)$$

Skyline search: non-dominated (Pareto) ranking

Drawback of scalarization: need fix λ; unknown user-dependent.

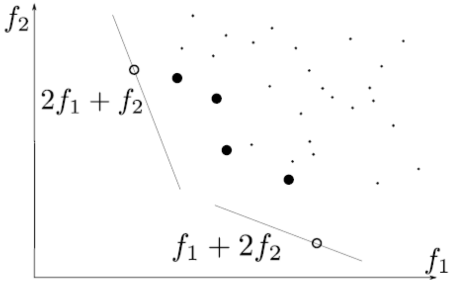

Non-dominated sorting: a point i is non-dominated if there exists no other point j such that $f_1(j) < f_2(i)$ and $f_2(j) < f_2(i)$.

The set of non-dominated points is an "antichain" called the Pareto front. Set of Pareto fronts is canonical antichain partition.

Papadias, Tao, Fu, Seeger, An Optimal and Progressive Algorithm for Skyline Queries, SIGMOD 2003

Searching for matches to dual queries over Pareto fronts

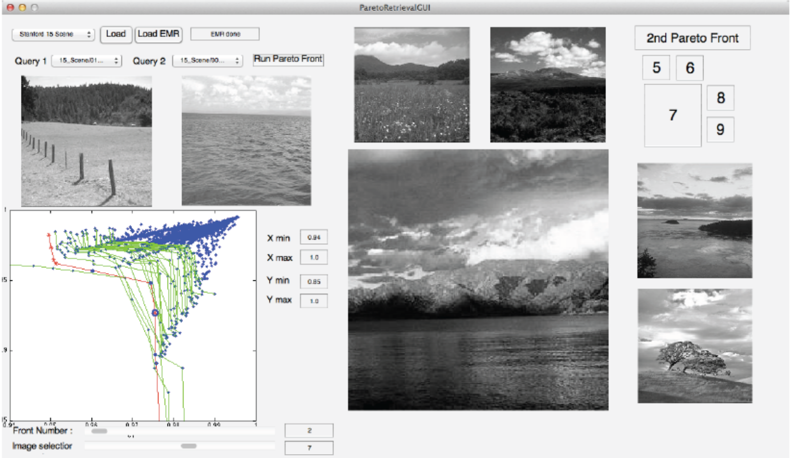

Hsiao, Calder and H, "Multiple-query Image Retrieval using Pareto Front Method," submitted.

Are there interesting mathematical issues?

- Is there an asymptotic theory (large n) for shape of the Pareto fronts T_k?
- What is the distribution of points over T_k?
- Can computational complexity of finding Pareto fronts be reduced?

Are there interesting mathematical issues?

- Is there an asymptotic theory (large p) for shape of the Pareto front T_k?
  - $\Rightarrow$ Yes. The Pareto front is solution to a pde on $\mathbb{R}^d$.
- What is the distribution of points on T_k?
  - $\Rightarrow$ $E[N_{Pareto}] = \gamma n^{(d-1)/d} + O(n^{(d-2)/d})$ with

$$\gamma = d^{-1}(d!)^{\frac{1}{d}}\Gamma(d^{-1}) \int_T f^{\frac{d-1}{2}}(u(z))(u_1(z)\cdots u_d(z))^{\frac{1}{d}}\, dz$$

- Can computational complexity of finding Pareto fronts be reduced?
  - $\Rightarrow$ Yes. "In principle" can reduce from $O(dn^2)$ to $O(1)$.

Asymptotic theorem

Let there be d criteria giving non-negative similarity scores $\mathbf{X}_i = [f_1(i), \ldots, f_d(i)]$ with ith image in the database, $i = 1, \ldots, n$.

Assume that $\{\mathbf{X}_i\}_{i=1}^n$ are i.i.d. from multivariate density $f(x_1, \ldots, x_d)$.

Theorem

As $n \to \infty$ the Pareto fronts converge uniformly to the level sets of the value function $U(x_1, \ldots, x_d)$ where U is the non-viscosity solution to the Hamilton-Jacobi partial differential equation:

$$\frac{\partial U}{\partial x_1} \cdot \ldots \cdot \frac{\partial U}{\partial x_d} = \frac{1}{d^d} f$$

Calder, Esedoglu and H, "A Hamilton-Jacobi equation for the continuum limit of non-dominated sorting," arXiv:1302.5828, 2013.

Illustration of Asymptotic Theory

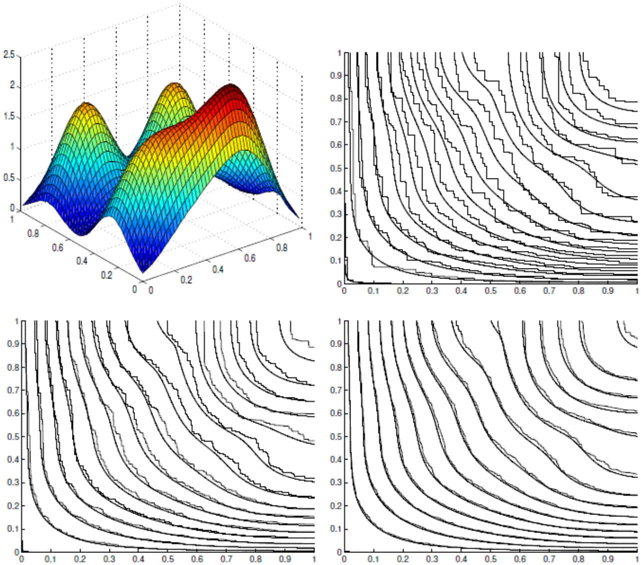

Outline

Some challenges

- Multimodality feature representation for processing/displaying multiple data types
- Tractable human performance modeling at suitable level of granularity
- Computational theory of man-machine interaction - what becomes computable with human assistance? What happens to learning rates?
- Competency acquisition - how rapidly can computer and human learn to best interact?
- Uncertainty quantification - feedback destroys i.i.d. property and standard tests of significance cannot be applied.

What fields are relevant to man-machine modeling?

- CS, Statistics and Applied Math: multi-agent systems, reinforcement learning, mathematics of perception
- Human factors research: pilot training, interactive displays, air traffic control
- Experimental psychology: cognitive psychometrics, perception, memory, cognition
- Health research: aging, stroke, cognitive impairments, psychological assessment
- Education research: early learning and cognitive development
- Competency assessment: ETS, computerized testing

Jeff Calder, Selim Esedoglu, and Alfred O Hero. A hamilton-jacobi equation for the continuum limit of non-dominated sorting. *arXiv preprint arXiv:1302.5828*, 2013.

Bryan R Gibson, Kwang-Sung Jun, and Xiaojin Zhu. With a little help from the computer: Hybrid human-machine systems on bandit problems. In *NIPS 2010 Workshop on Computational Social Science and the Wisdom of Crowds*, 2010.

Ko-Jen Hsiao, Kevin S Xu, Jeff Calder, and Alfred O Hero III. Multiple-query image retrieval using pareto front method. In *Proc. of Advances in Neural Information Processing Systems (NIPS)*, 2013.

Kevin G Jamieson, Robert D Nowak, and Benjamin Recht. Query complexity of derivative-free optimization. *arXiv preprint arXiv:1209.2434*, 2012.

Bruno Jedynak, Peter I Frazier, and Raphael Sznitman. Twenty questions with noise: Bayes optimal policies for entropy loss. *Journal of Applied Probability*, 49(1):114–136, 2012.

Ki Hang Kim and Fred William Roush. *Team theory*. E. Horwood, 1987.

K Ronald Laughery, Christian Lebiere, and Susan Archer. Modeling human performance in complex systems. *Handbook of human factors and ergonomics*, pages 965–996, 2006.

Mark Steyvers, Michael D Lee, and Eric-Jan Wagenmakers. A bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53(3):168–179, 2009.

T. Tsiligkaridis and A.O. Hero. A collaborative 20 questions model for target search with human-machine interaction. *arXiv:1302.5828*, June 2013.

E.-J. Wagenmakers, L.J. Maas, and P.P. Grasman. An ez-diffusion model for response time and accuracy. *Psychonomic Bulletin and Review*, 14(1):3–22, 2007. ISSN 1069-9384. doi: 10.3758/BF03194023. URL http://dx.doi.org/10.3758/BF03194023.

X. Wang, J. O. Berger, and D. S. Burdick. Bayesian analysis of dynamic item response models in educational testing. *ArXiv e-prints*, April 2013.

Two views of human evolution