"Medicine and the Computer: The Promise and Problems of Change"

—W.B. Schwartz, NEJM 1970

- Perceived problems
  - Physician shortage and maldistribution
  - Ever-expanding body of knowledge, so that the physician cannot keep up
- Exploit the computer as an "intellectual", "deductive" instrument
  - Improve medical care
  - Separate practice from memorization
  - Allow time for human contact
  - Encourage different personalities in medicine — the "healing arts"

Tasks?

- Diagnosis
- Prognosis
- Therapy

"One-shot"
vs.
Ongoing

- **"Doctor's Assistant" for clinicians at any level of training**
- **Expert (specialist) consultation for non-specialists**
- **Monitoring and error detection**
- **Critiquing, what-if**
- **Guiding patient-controlled care**
- **Education and Training**
- **Contribution to medical research**
- **…**

Two Historical Views on How to Build Expert Systems

- Great cleverness
  - Powerful inference abilities
  - *Ab initio* reasoning
- Great stores of knowledge
  - Possibly limited ability to infer, but
  - Vast storehouse of relevant knowledge, indexed in an easy-to-apply form

How to do diagnosis (medical reasoning)?

- Program it
  - use a flowchart (since 1950's)
  - use rules (since 1970's)
- Deduce it
  - use some representation of disease and a diagnostic algorithm
    - disease/symptom associations (since 1960's)
    - probabilistic version (since 1960's)
    - causal models (since 1980's)

Flowcharts contain all of...

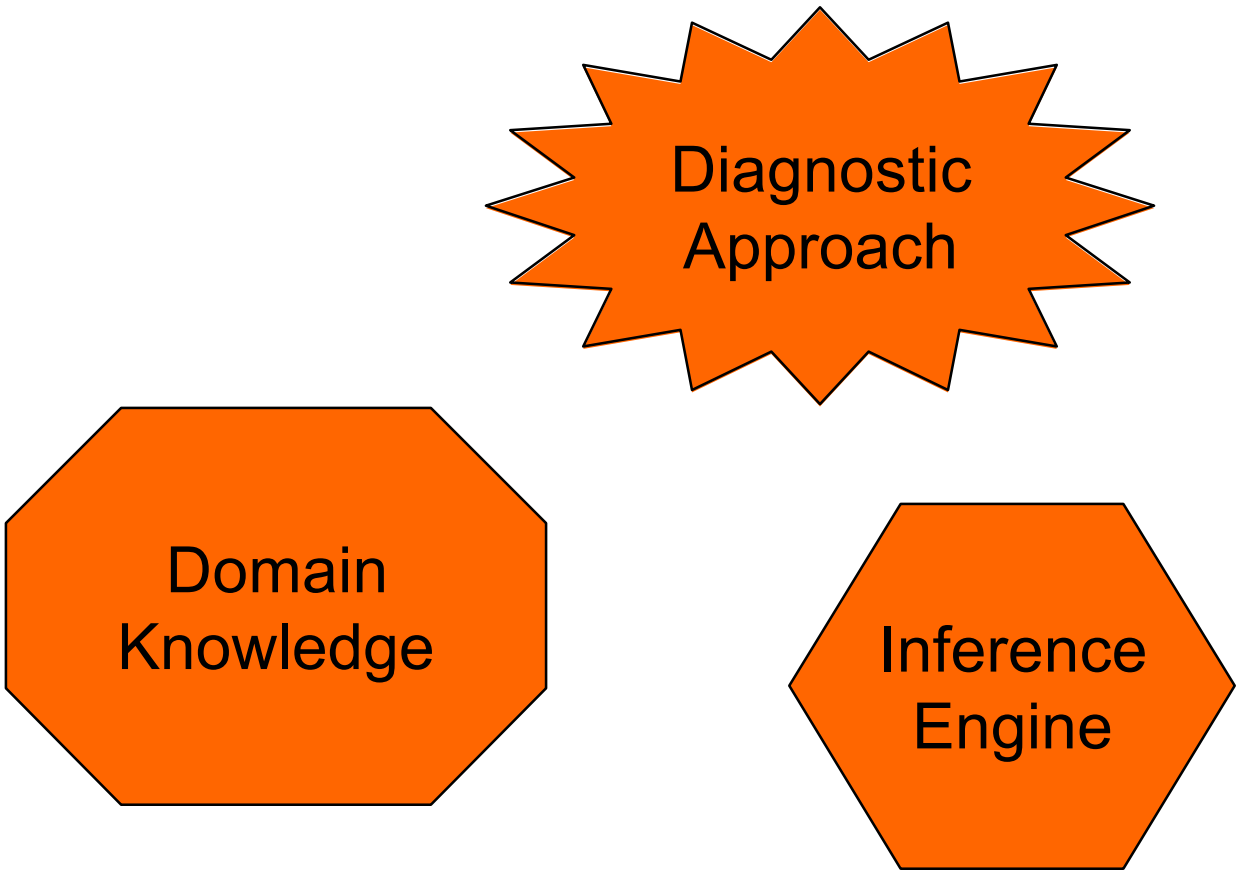

Flowcharts

- Good:
  - Simple
  - Easy to build
- Bad:
  - Hard to deal with
    - missing data
    - out of sequence data
    - uncertainty
  - Hard to maintain

Mycin—Rule-based Systems

- **Task: Diagnosis and prescription for bacterial infections of the blood (and later meningitis)**

- **Method:**
  - **Collection of modular rules**
  - **Backward chaining**
  - **Certainty factors**

RULE037

IF the organism
1) stains grampos
2) has coccus shape
3) grows in chains

THEN
There is suggestive evidence (.7) that the identity of the organism is streptococcus.

Mycin consult

Davis, *et al.*, *Artificial Intelligence* **8**: 15-45 (1977)

How Mycin Works

- To find out a fact
  - If there are rules that can conclude it, try them
  - Ask the user
- To “run” a rule
  - Try to find out if the facts in the premises are true
  - If they all are, then assert the conclusion(s), with a suitable certainty
- Backward chaining from goal to given facts

- Dynamically traces out *behavior* of (what might be) a flowchart
  - Information used everywhere appropriate
  - Single expression of any piece of knowledge

Explore Mycin's Use of Knowledge

** **Did you use RULE 163 to find out anything about ORGANISM-1?**

RULE163 was tried in the context of ORGANISM-1, but it failed because it is not true that the patient has had a genito-urinary tract manipulative procedure (clause 3).

** **Why didn't you consider streptococcus as a possibility?**

The following rule could have been used to determine that the identity of ORGANISM-1 was streptococcus: RULE033

But clause 2 ("the morphology of the organism is coccus") was already known to be false for ORGANISM-1, so the rule was never tried.

Davis, *et al.*, *Artificial Intelligence* **8**: 15-45 (1977)

Mycin contains ...

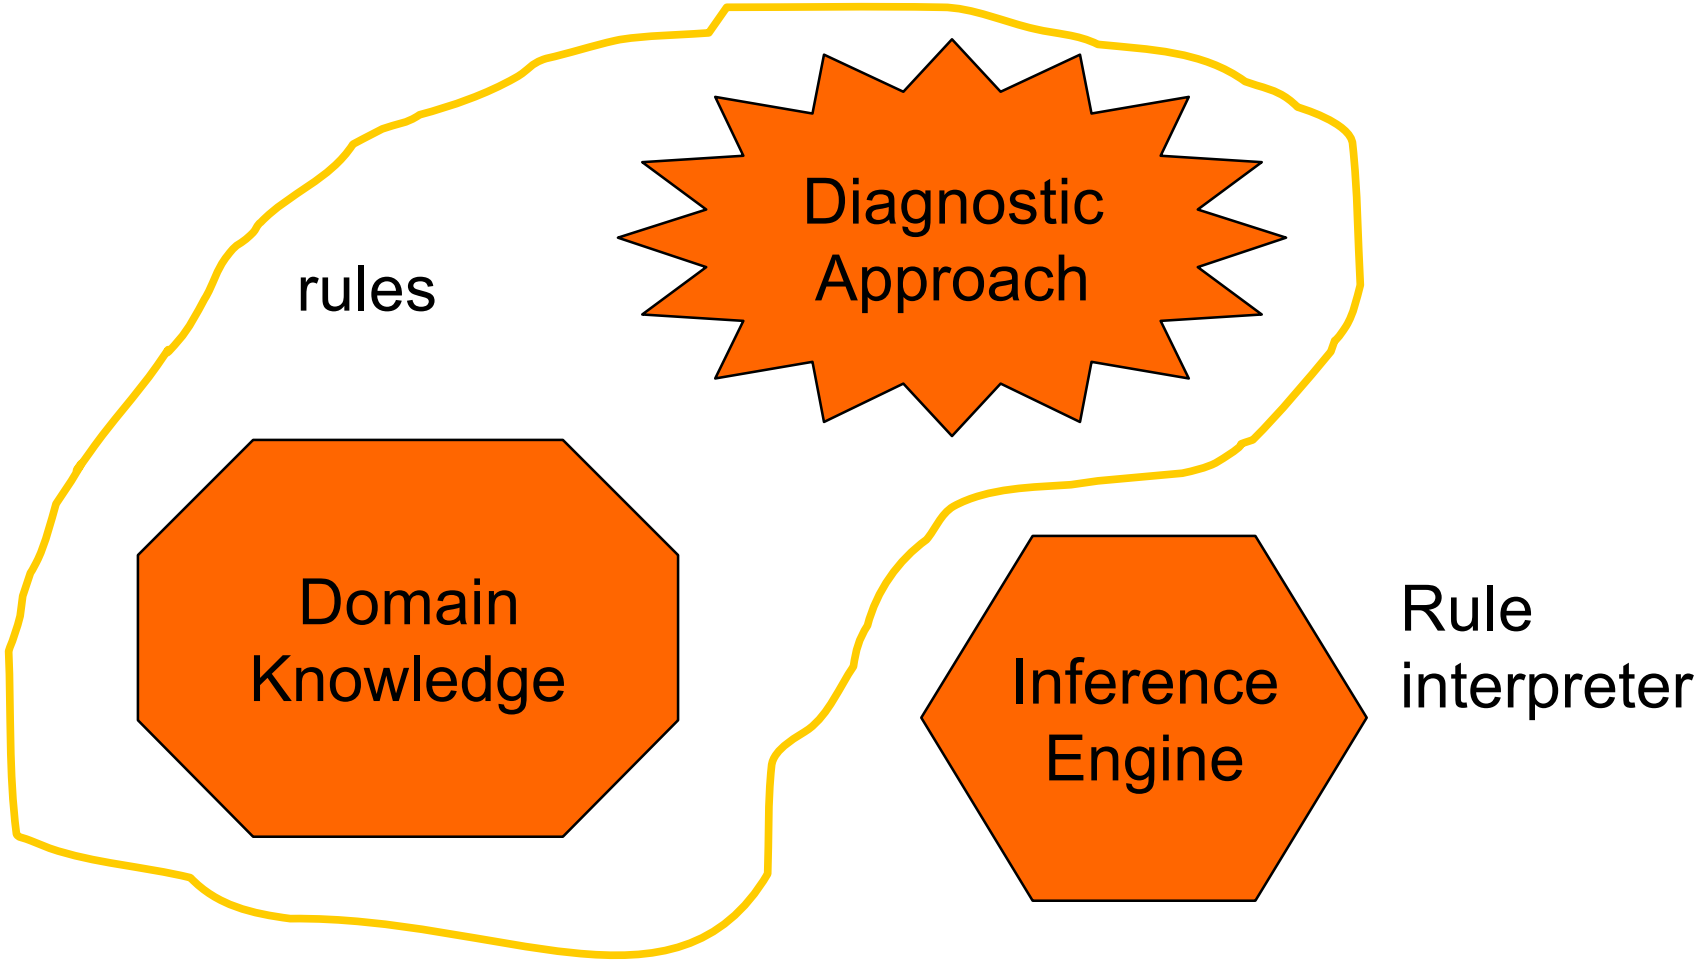

Representation

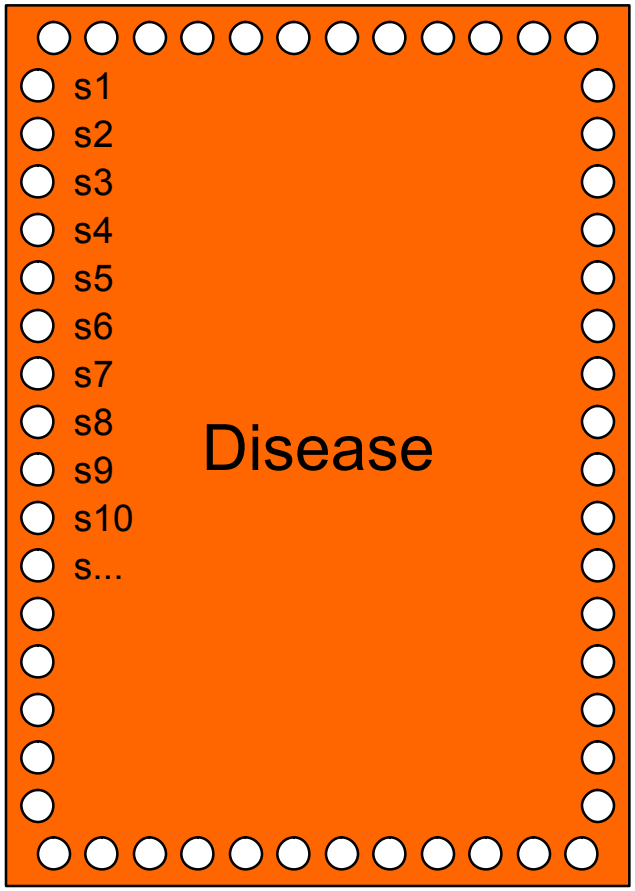

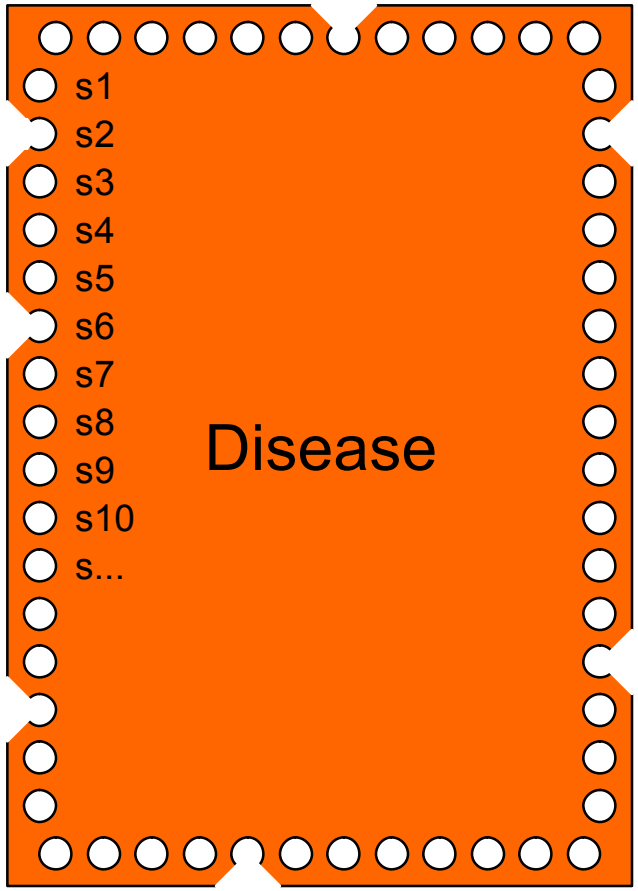

Diagnosis by Card Selection

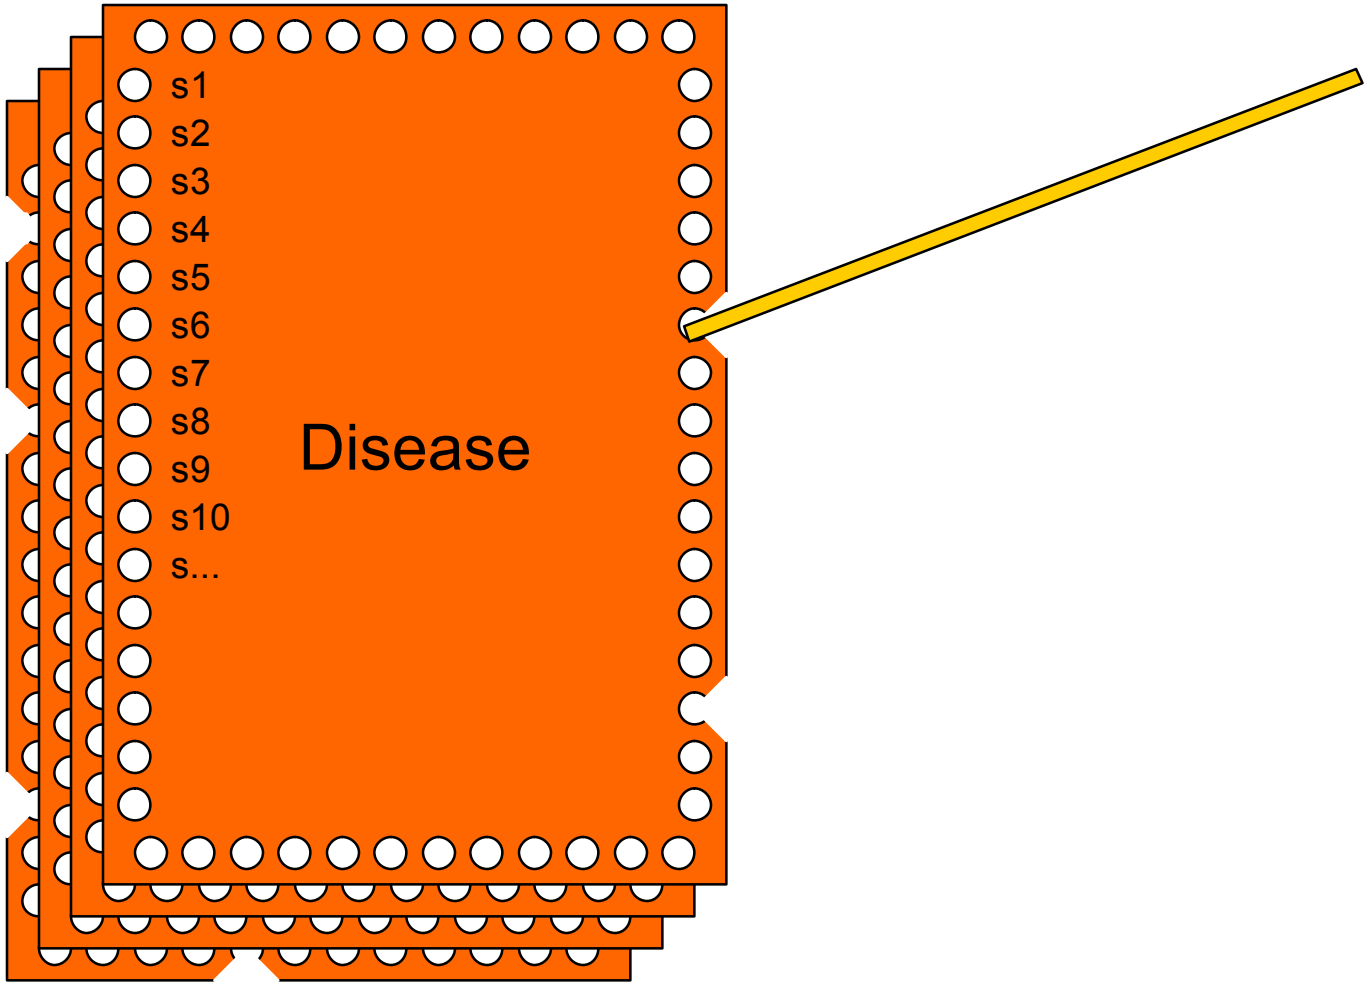

Diagnosis by Edge-Punched Cards

- Dx is intersection of sets of diseases that *may cause* all the observed symptoms
- Difficulties:
  - Uncertainty
  - Multiple diseases

~ “Problem-Knowledge Coupler” of Weed

Probabilistic Version of Cards

- Assume single disease
- Symptoms depend only on disease state
  - Conditional independence
    - P(s,t|d) = P(s|d)P(t|d)
- Bayes' Rule updates disease probabilities based on observing symptoms
- Next lecture's large example

Taking the Present Illness—Diagnosis by Pattern Directed Matching

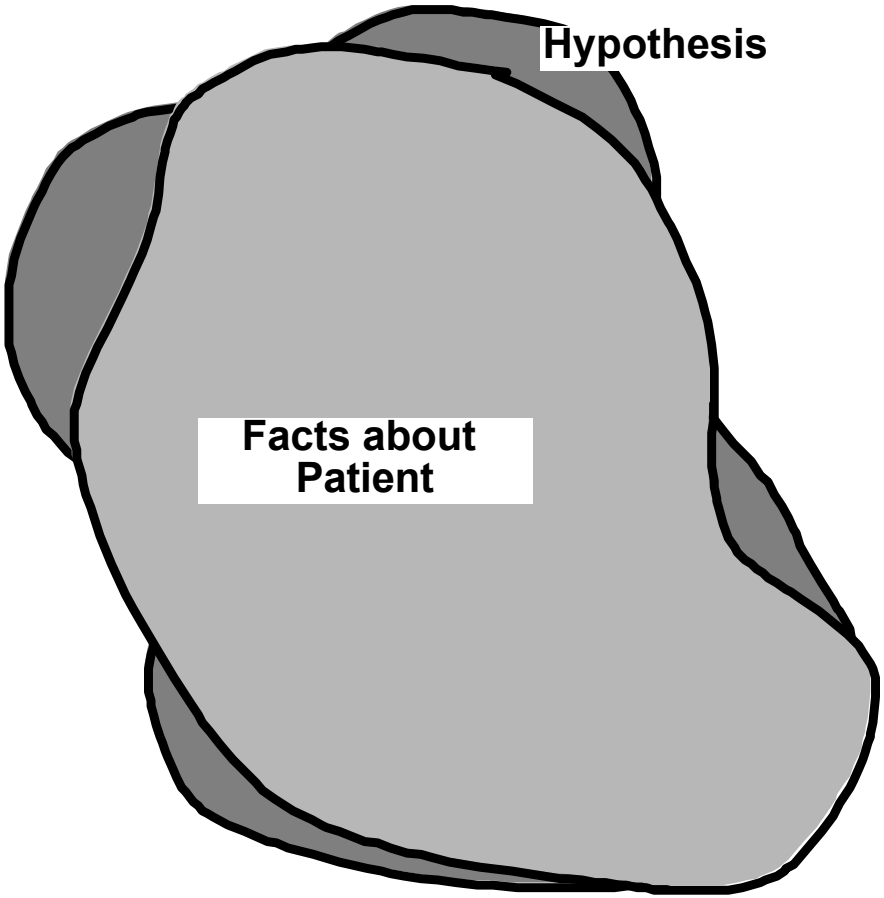

PIP's Theory of Diagnosis

- **From initial complaints, *guess* suitable hypothesis.**
- **Use current active hypotheses to guide questioning**
- **Failure to satisfy expectations is the strongest clue to a better hypothesis; *differential diagnosis***
- **Hypotheses are *activated, de-activated, confirmed* or *rejected* based on**
  (1) logical criteria
  (2) probabilities based on:
  findings local to hypothesis
  causal relations to other hypotheses

Memory Structure in PIP

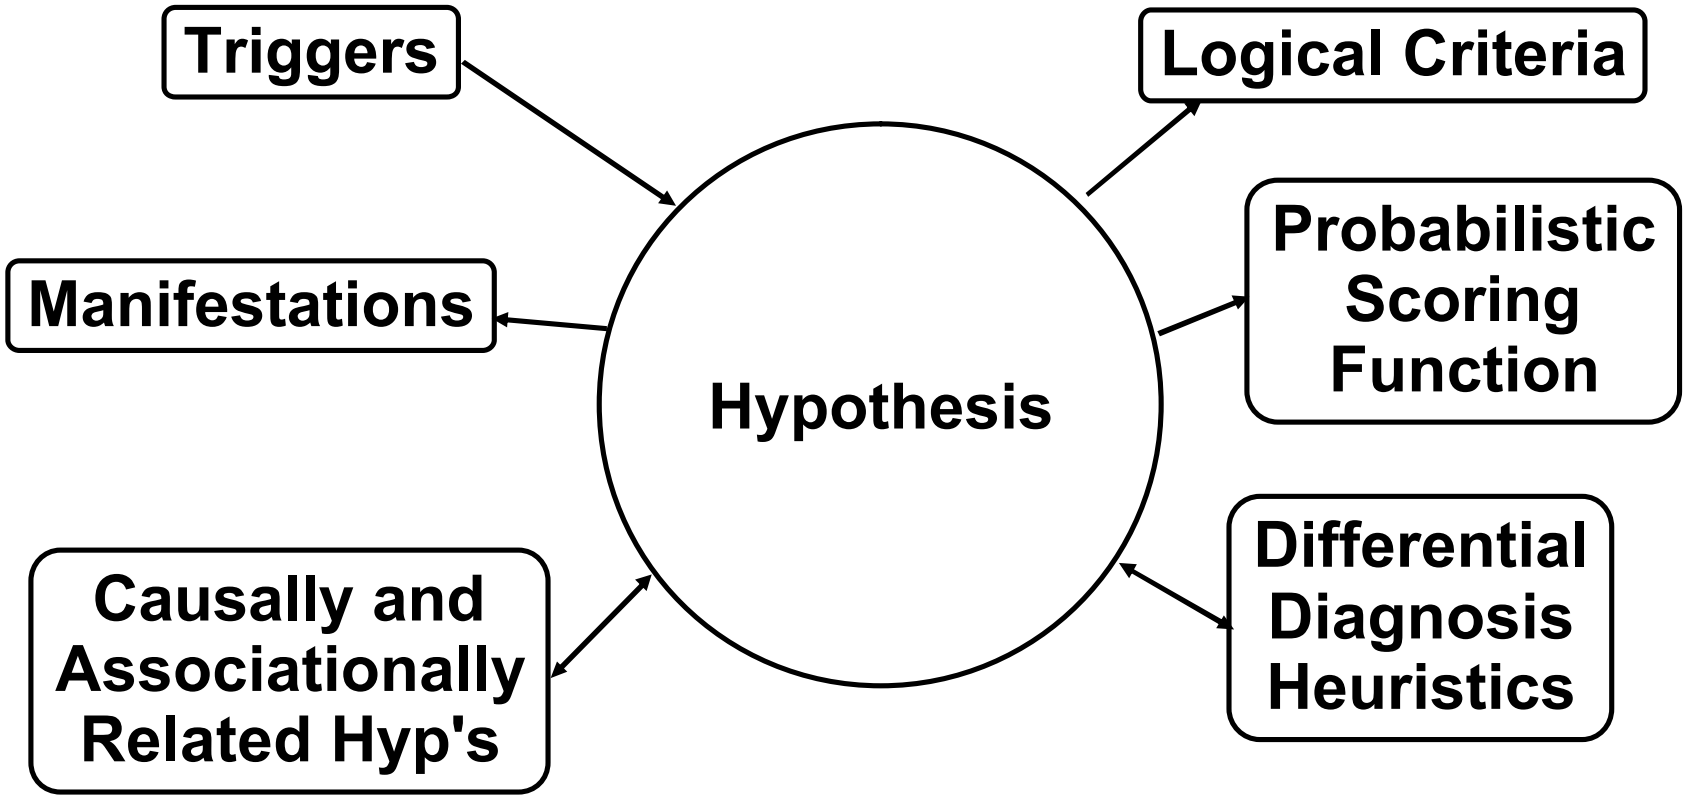

PIP's Model of Nephrotic Syndrome

NEPHROTIC SYNDROME, *a clinical state*

FINDINGS:

1* Low serum albumin concentration
2. Heavy proteinuria
3* >5 gm/day proteinuria
4* Massive symmetrical edema
5* Facial or peri-orbital symmetric edema
6. High serum cholesterol
7. Urine lipids present

IS-SUFFICIENT: **Massive pedal edema & >5 gm/day proteinuria**

MUST-NOT-HAVE: **Proteinuria absent**

SCORING . . .

MAY-BE-CAUSED-BY: **AGN, CGN, nephrotoxic drugs, insect bite, idiopathic nephrotic syndrome, lupus, diabetes mellitus**

MAY-BE-COMPLICATED-BY: **hypovolemia, cellulitis**

MAY-BE-CAUSE-OF: **sodium retention**

DIFFERENTIAL DIAGNOSIS:

neck veins elevated → constrictive pericarditis
ascites present → cirrhosis
pulmonary emboli present → renal vein thrombosis

PIP's Analysis of a Case

PRESENTING SYMPTOMS: EDEMA, ERYTHEMATOUS, PITTING, SYMMETRICAL, WORSE-IN-EVENING, FIRST-TIME, FOR-DAYS AND MASSIVE. HE DOES NOT HAVE DYSPNEA. HE HAS SOCIAL ALCOHOL CONSUMPTION. HE DOES NOT HAVE JAUNDICE. IT IS NOT EXPLICITLY KNOWN WHETHER IN THE PAST HE HAD PROTEINURIA, BUT HE HAS SMALL-POLICY LIFE INSURANCE, AND HE HAS SERVED-IN ARMED FORCES. HE DOES NOT HAVE VARICOSE VEINS. IN THE PAST HE DID NOT HAVE EDEMA. HE DOES NOT HAVE HEMATURIA. HE HAS NORMAL BUN. HE HAS NORMAL CREATININE. HE HAS PERI-ORBITAL EDEMA, WHICH IS WORSE-IN-MORNING, FIRST-TIME, FOR-DAYS AND SYMMETRICAL. …

DIAGNOSES THAT HAVE BEEN ACCEPTED ARE: NEPHROTIC SYNDROME AND SODIUM RETENTION.

THE LEADING HYPOTHESIS IS IDIOPATHIC NEPHROTIC SYNDROME.

	fit	*explained*	*score*
IDIOPATHIC NEPHROTIC SYNDROME	0.80	0.37	0.58
ACUTE GLOMERULONEPHRITIS	0.22	0.27	0.24
HENOCH SCHOENLEIN PURPURA	0.07	0.10	0.09

Other "Frame-based" Systems

- Internist/QMR
- DXPLAIN
- ILIAD
  - Local Bayesian models

Internist/QMR

- Knowledge Base:
  - 956 hypotheses
  - 4090 manifestations (about 75/hypothesis)
  - *Evocation* like P(H|M)
  - *Frequency* like P(M|H)
  - *Importance* of each M
  - *Causal relations* between H's
- Diagnostic Strategy:
  - Scoring function
  - Partitioning
  - Several questioning strategies

QMR Scoring

- Positive Factors
  - Evoking strength of observed Manifestations
  - Scaled Frequency of causal links from confirmed Hypotheses
- Negative Factors
  - Frequency of predicted but absent Manifestations
  - Importance of unexplained Manifestations
- Various scaling parameters (roughly exponential)

QMR Partitioning

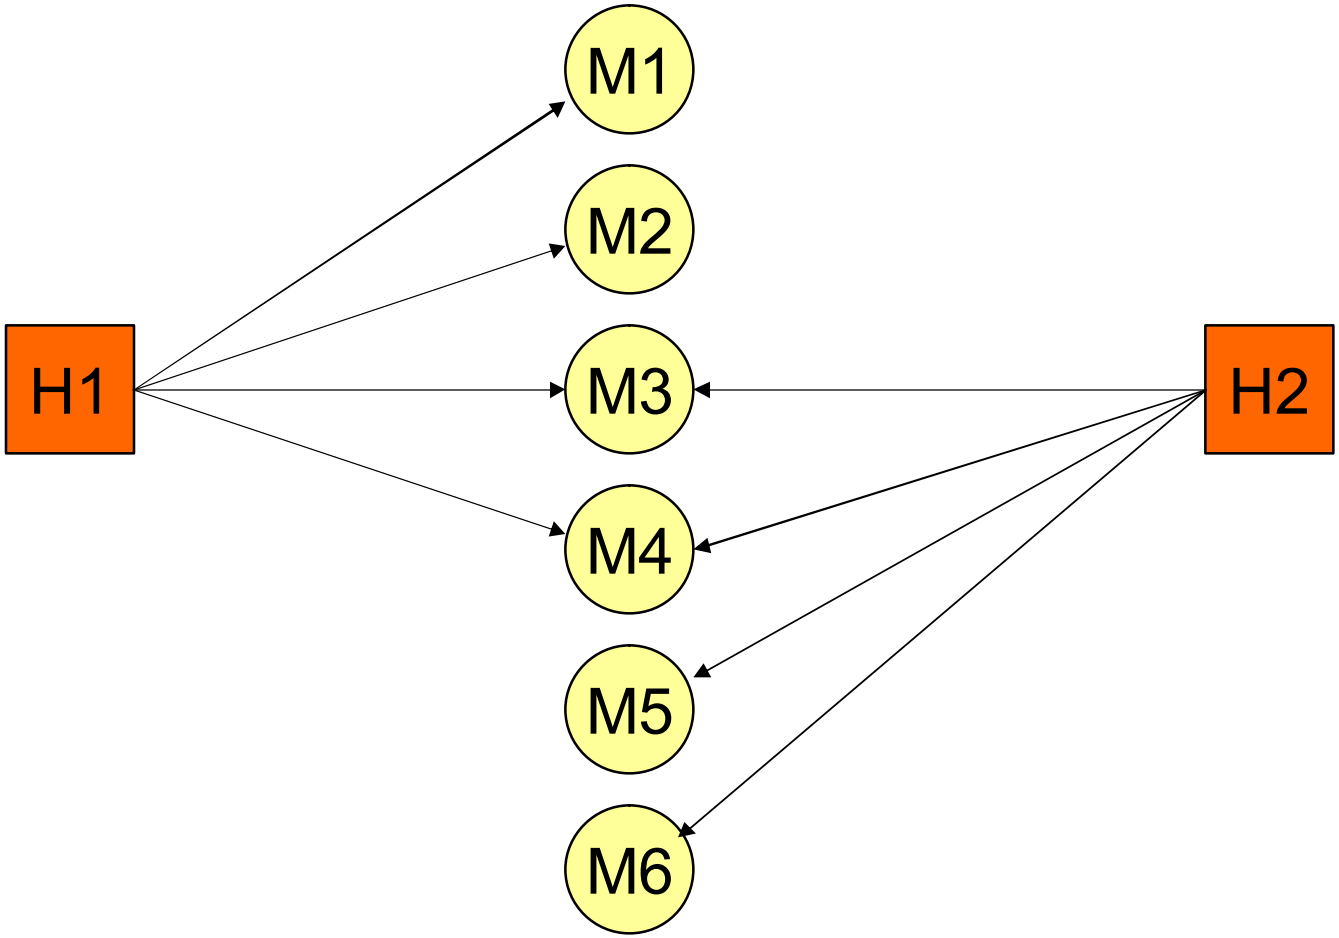

Competitors

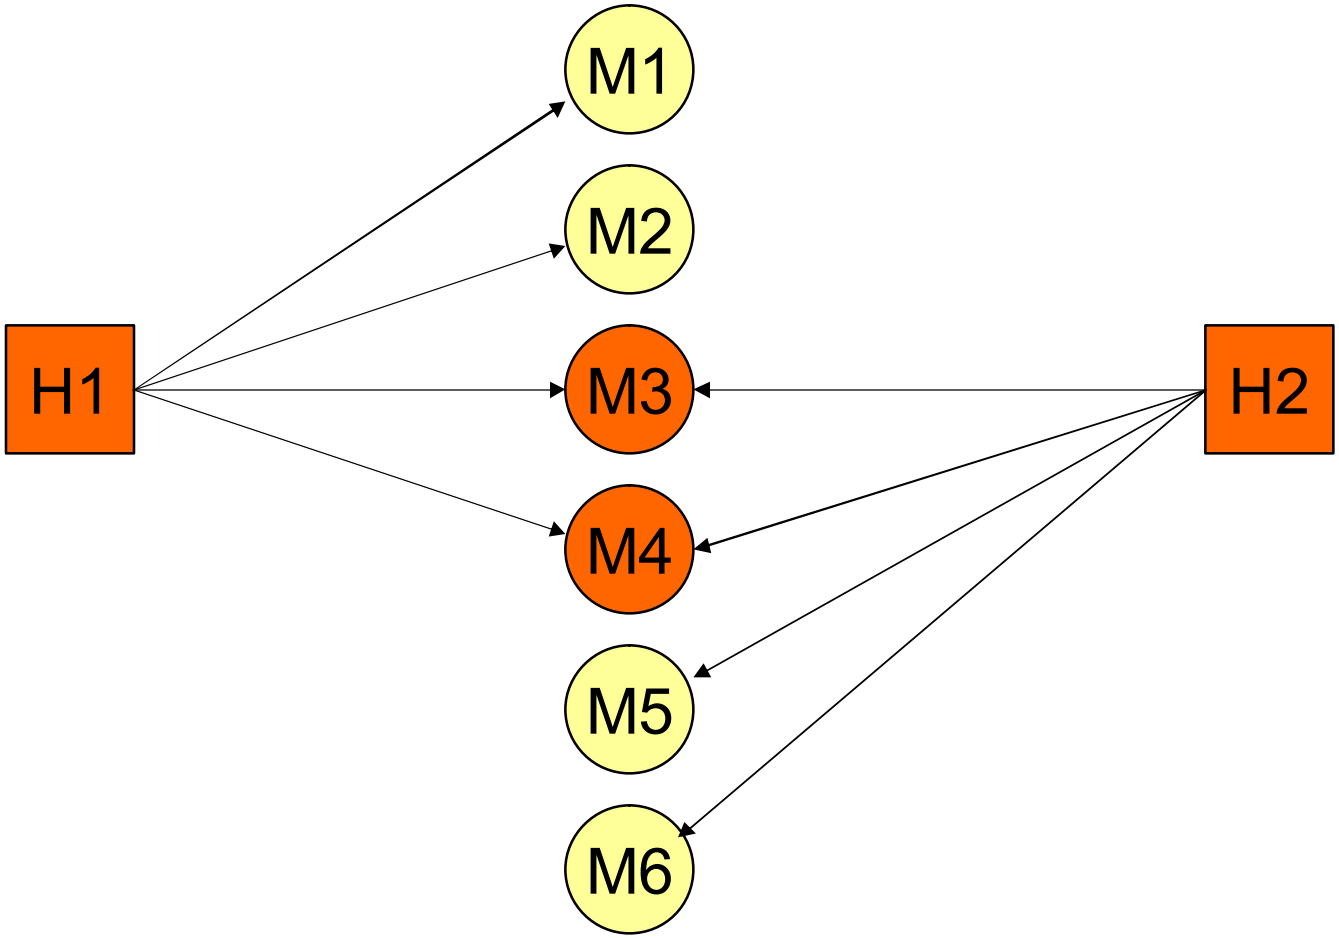

Still Competitors

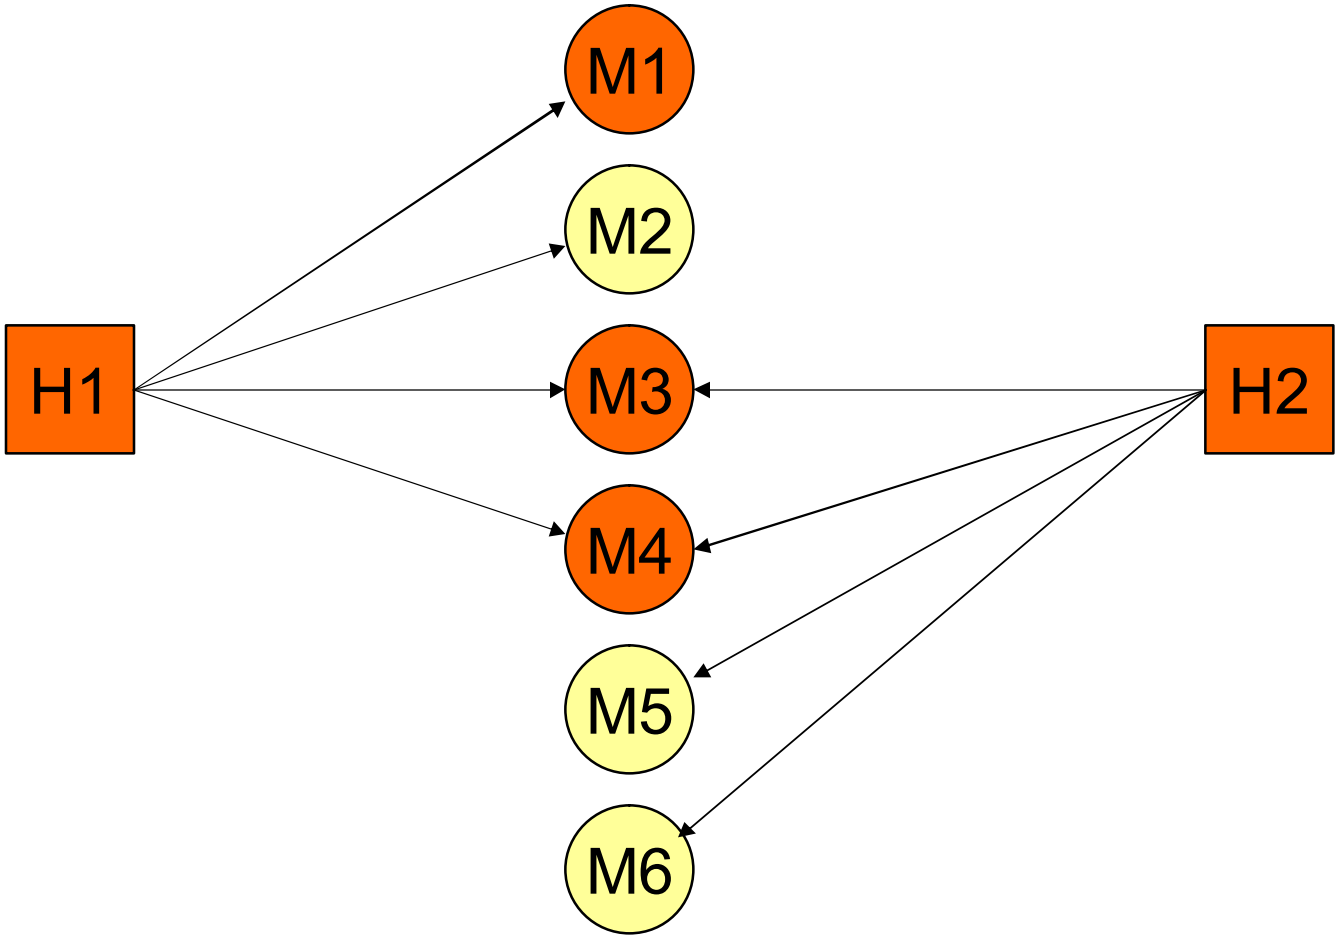

Probably Complementary

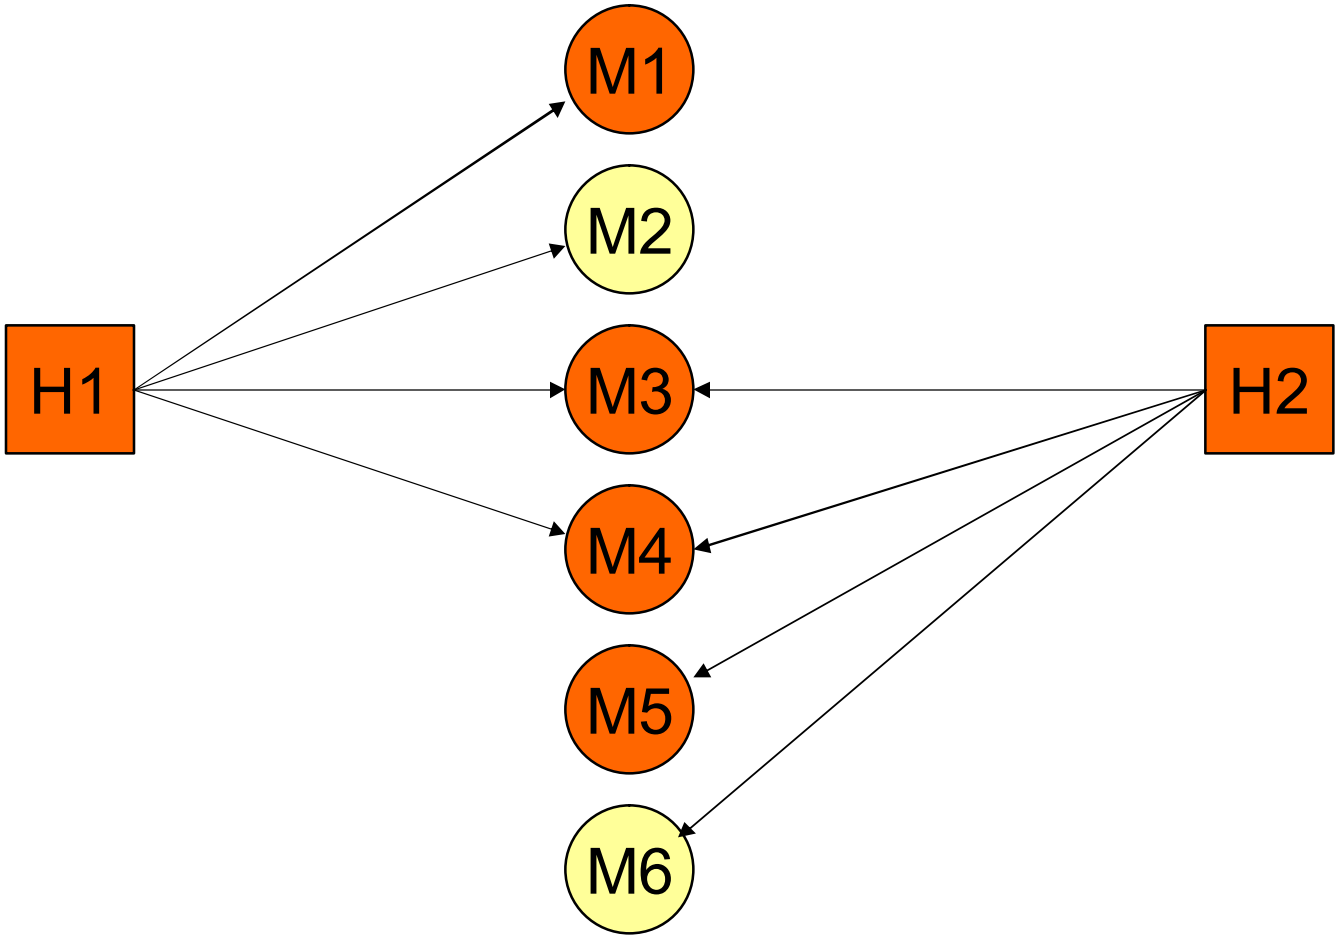

Multi-Hypothesis Diagnosis

- Set aside complementary hypotheses
- … and manifestations predicted by them
- Solve diagnostic problem among competitors
- Eliminate confirmed hypotheses and manifestations explained by them
- Repeat as long as there are coherent problems among the remaining data

Frame-based Diagnosis

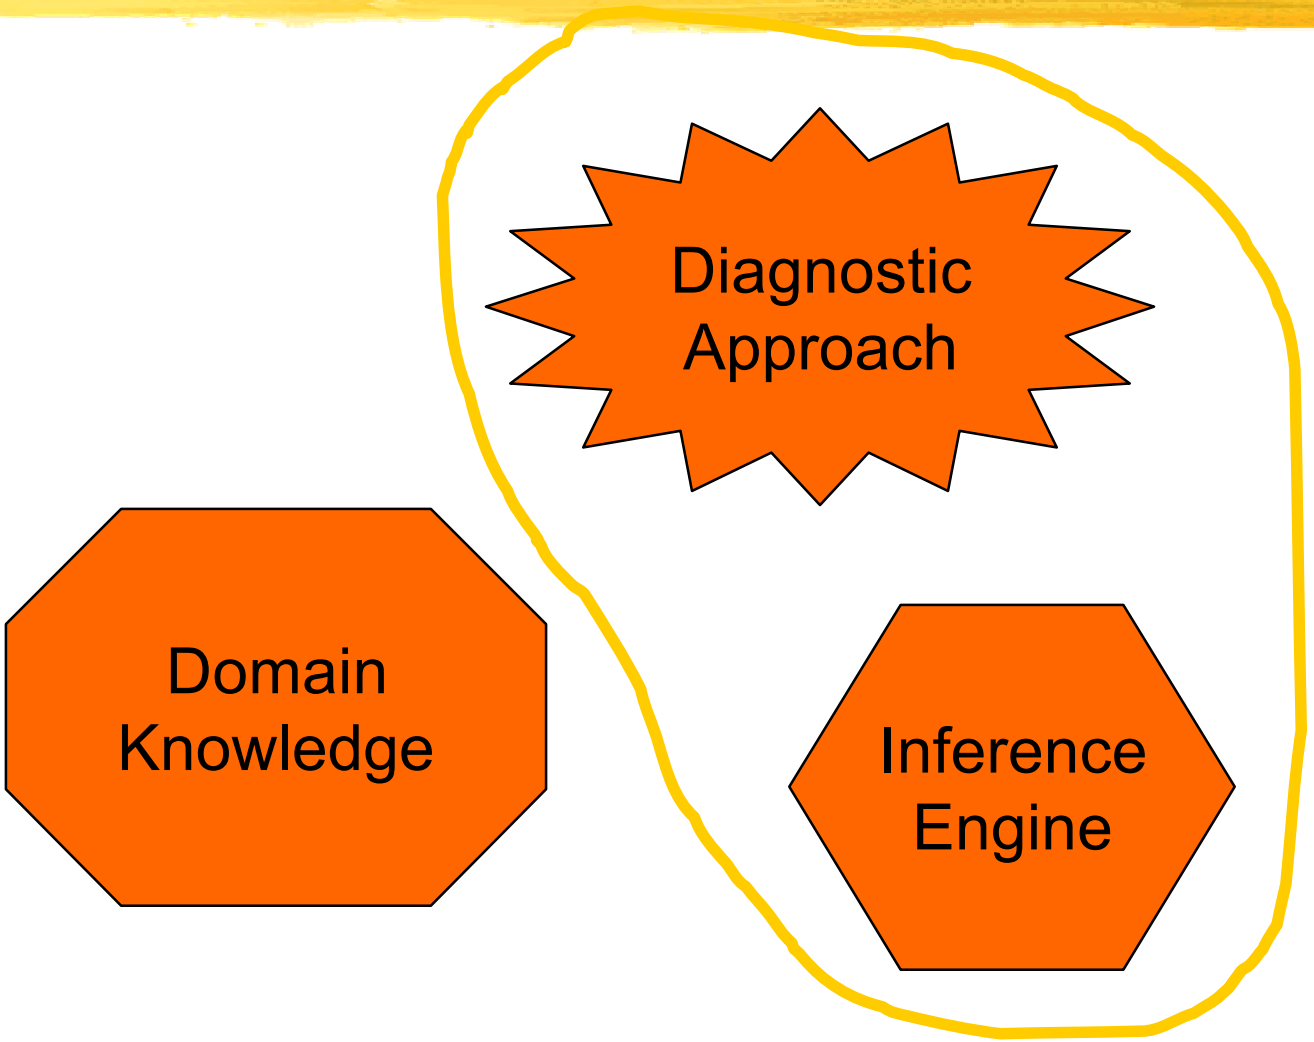

Problems with Dx Programs

- *Wonderful* for very limited domain, but for general medicine:
  - Not very accurate
  - Very difficult to build & maintain
  - Unsophisticated reasoning
    - time
    - space
    - severity
    - causality
  - Little exploitation of data

What do People Know?

- Human expertise appears to be more than statistical association
- Medical knowledge:
  - physiology
  - pathophysiology
  - pathology
  - genetics, . . .
- Clinical knowledge:
  - focus of attention
  - following a process
  - heuristics

The Surprisingly Normal pH

- Diarrhea causes bicarbonate (alkali) loss
- Vomiting causes acid loss
- Therefore, normal pH is a manifestation of {diarrhea + vomiting}!

Temporal Reasoning

- **Keeping track of multiple forms of temporal relations (Kahn '75)**
  - **The time line**
    - **"On Dec. 12 last year . . ."**
  - **Special reference events**
    - **"Three days after I was hospitalized in 1965 . . ."**
  - **Temporal Ordering Chains**
    - **"It must have been before I graduated from high school."**

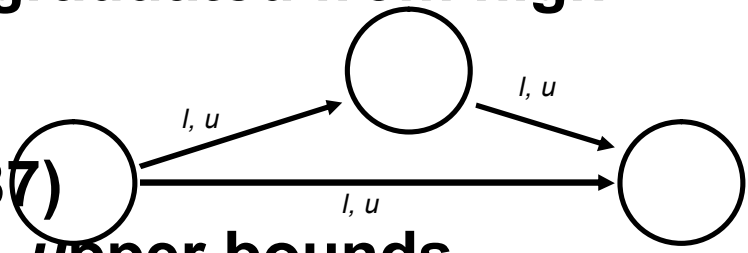

- **Constraint propagation (Kohane '87)**
  - **Primitive relation: e1, e2, *l*ower, *u*pper bounds**
  - **Heuristics for propagation based on semantic grouping**

Exploiting Temporal Relations

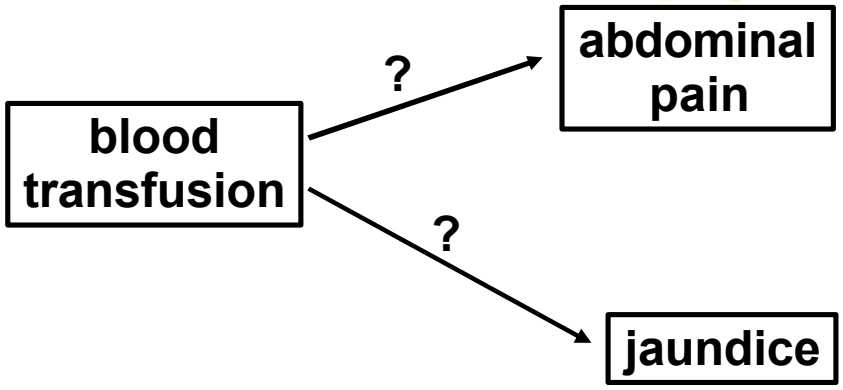

- transfusion precedes both abdominal pain and jaundice *implies* transfusion-borne acute hepatitis B
- as in 1, but only by one day
- jaundice occurred 20 years ago, transfusion and pain recent
- Can be very efficient at filtering out nonsense hypotheses.

Interpreting the Past with a Causal/Temporal Model

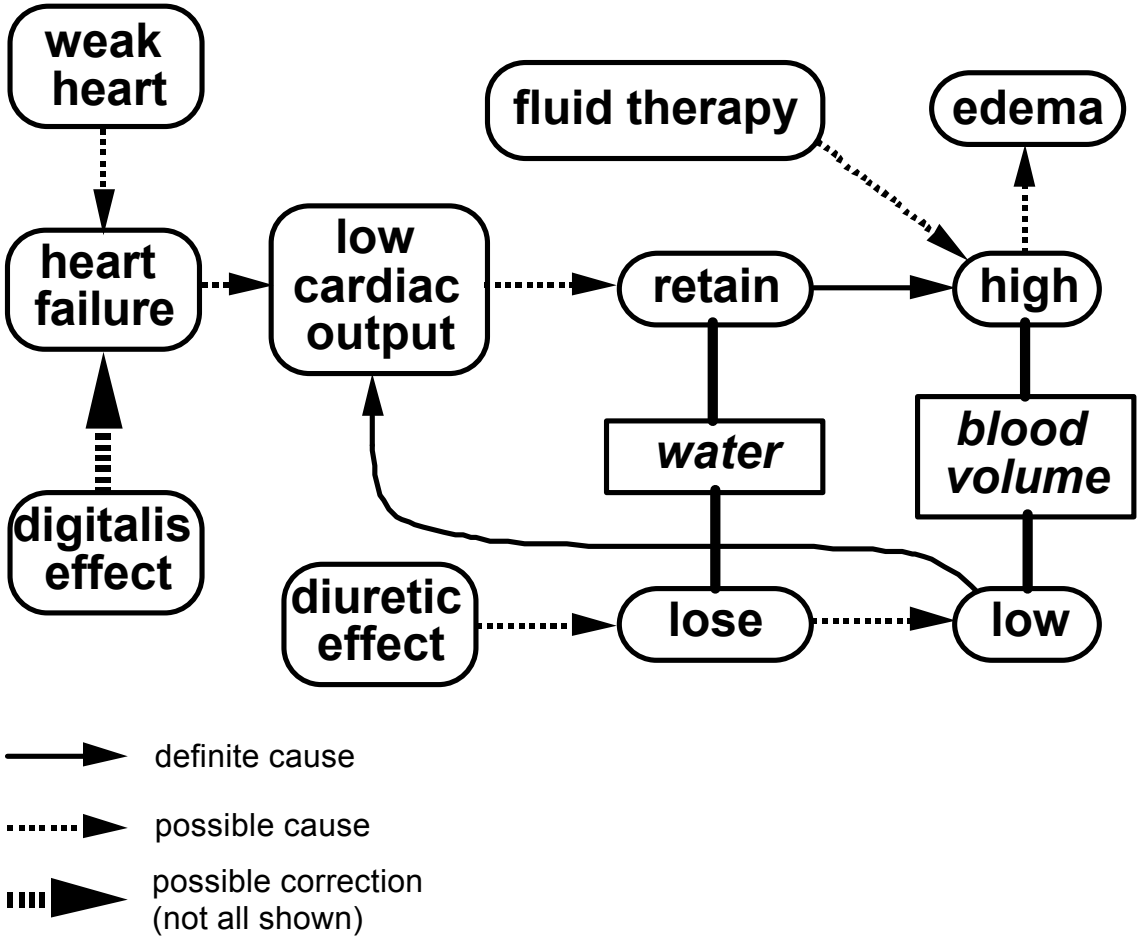

Postdiction

Long, Reasoning about State from Causation and Time in a Medical Domain, *AAAI 83*

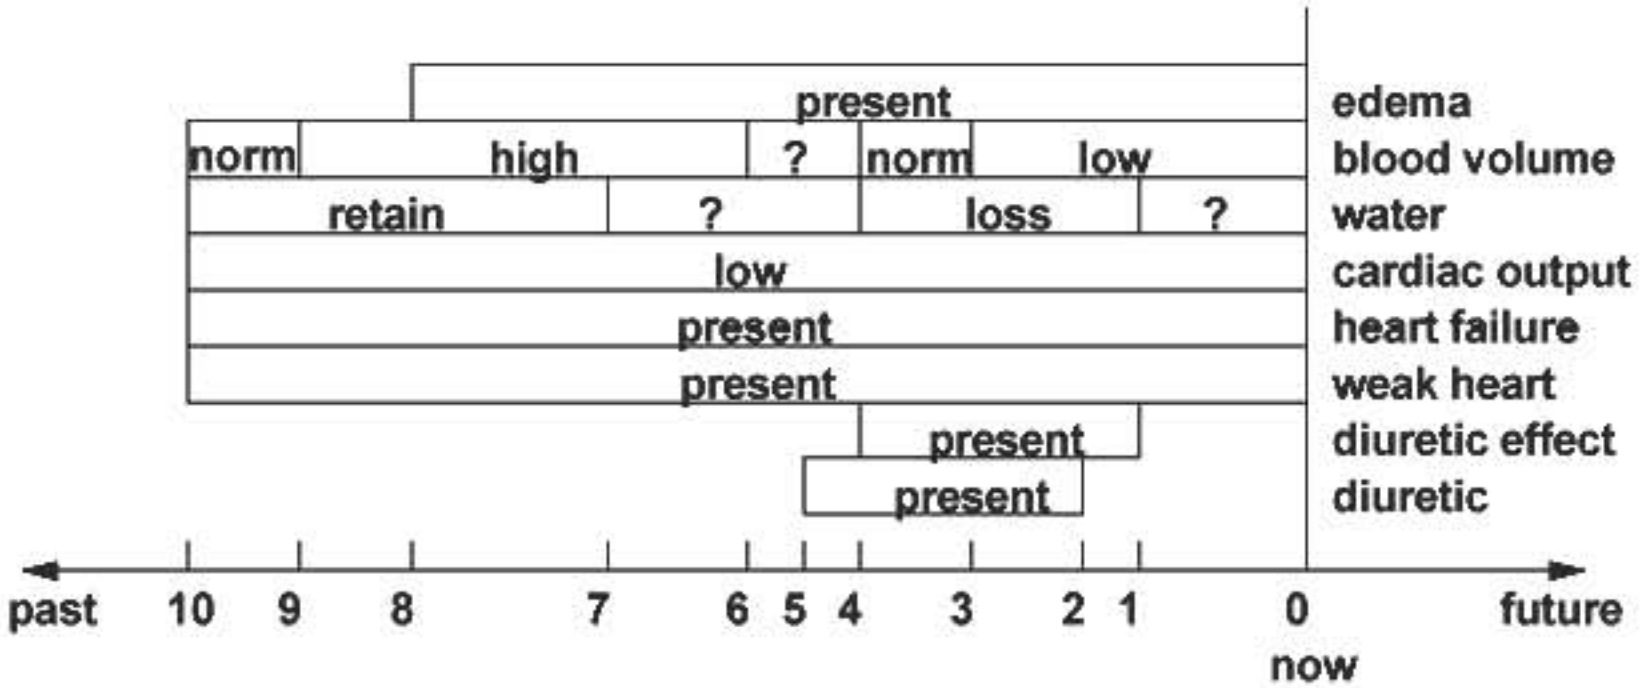

Reasoning from Models

- Model handles all possible interactions, without having explicitly to anticipate them all
- Strawman: Fit parameters to a physiological model, then predict consequences to suggest
  - other expected findings
  - reasonable interventions
- Qualitative models
- Combining associational and model-based reasoning

Guyton's Model of Cardiovascular Dynamics

Long's Clinical Model of Heart Failure Predictions for Mitral Stenosis with Exercise

Heart Disease Model

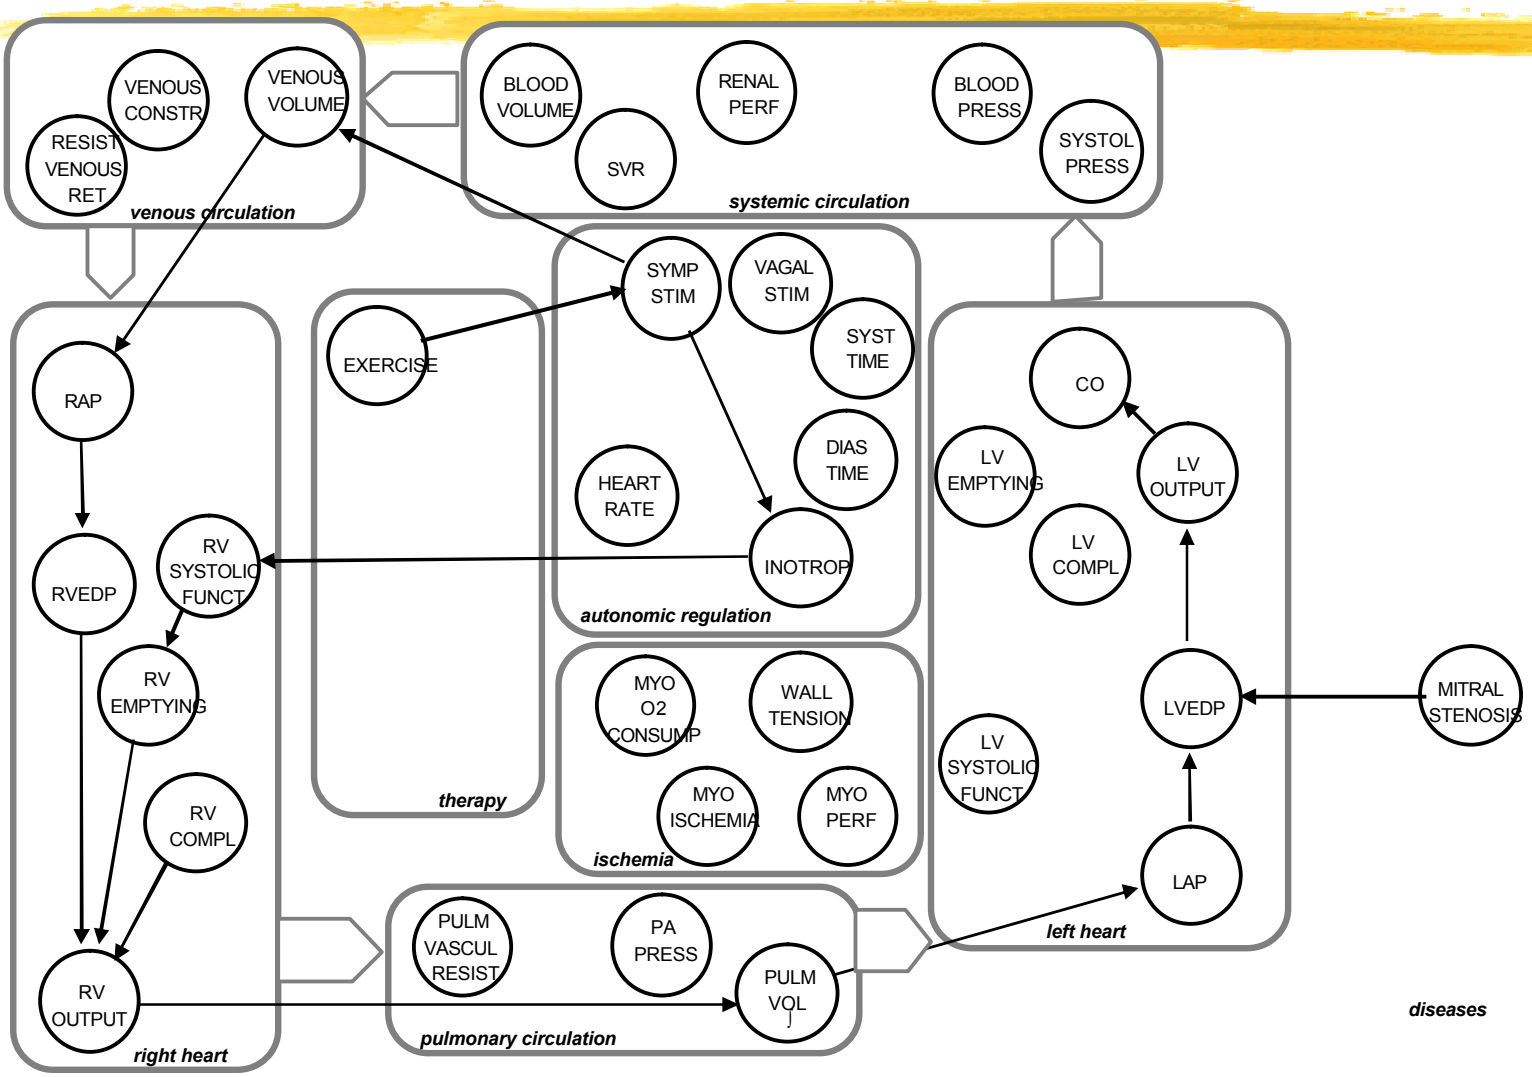

Multi-Level Causal Model

State of the Art (1989)

- Small, self-contained systems should be easy, but there are not very many being built.
  - • By contrast, Feigenbaum et al. point to 1,500 commercial systems in use in 1988, with thousands more in development
- A few sophisticated, modern, probability-based systems are now being built
- HIS's really are being developed (slowly, but surely) and will provide a critical opportunity for experimentation
- No large, broad-domain, deep systems are being tackled today
- Research advances are laying the groundwork for doing so in the future

State of Practice (today)

- Low-hanging fruit (important & tastes good)
  - "one-rule" expert systems
  - data presentation
- Knowledge ➔ Data
  - Classification, regression, neural networks, rough sets, fuzzy logic, Bayes nets, …
- Integration into clinical workflow
  - guidelines, care plans, …

6.872/HST951

- Emphasis is on learning from data
  - Thus, applied machine learning, various methods
  - Issues of data quality, evaluation of models
- Sensitivity to special needs of medical application