A Cognitive Routing System for the Internet

Dimitri Papadimitriou, & Benoit Donnet

8th Würzburg Workshop on IP:
Joint EuroNF, ITC, and ITG Workshop on
"**Visions of Future Generation Networks**"
July 21-22, 2008

Alcatel·Lucent

Outline

- Introduction
- Cognitive Routing
- Knowledge Plane
- Positioning
- Challenges and Perspectives
- Conclusion

Introduction

- Internet challenges and its evolution
  - Operational challenges: manageability/diagnosability, and availability
  - New challenges: security, accountability, and scalability (routing)
- Cognitive routing system: networking x machine learning technique

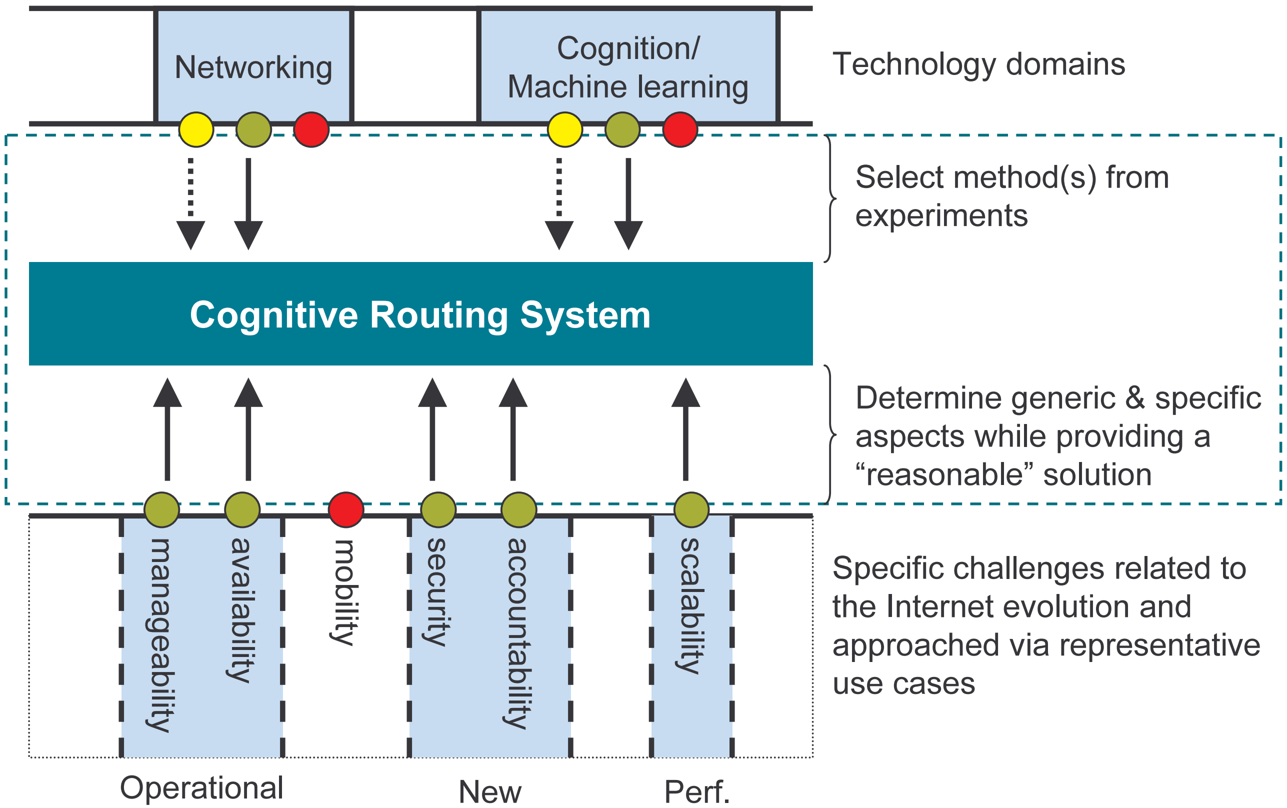

Why **learning** paradigm ?

- Similarity to the conditions traditionally encountered in classical machine learning problems:
  - **Nature**: events cannot be well characterized even when examples of such an event are available (inherent complexity in precisely characterizing an event)
  - **Relationship**: correlations and trends between events are hidden within large amounts of data associated to these events
  - **Environment**: changing conditions over time (in part., for routing system but also variability of user demands, expectations & behaviours)
  - **Quantity**: amount of available data is too large for handling by human intervention
  - **Evolutive**: new events are constantly detected/discovered
- Main concept: extend IP networking equipment, with a distributed cognitive engine based on semi-supervised, on-line, and distributed machine learning techniques

Cognitive Routing

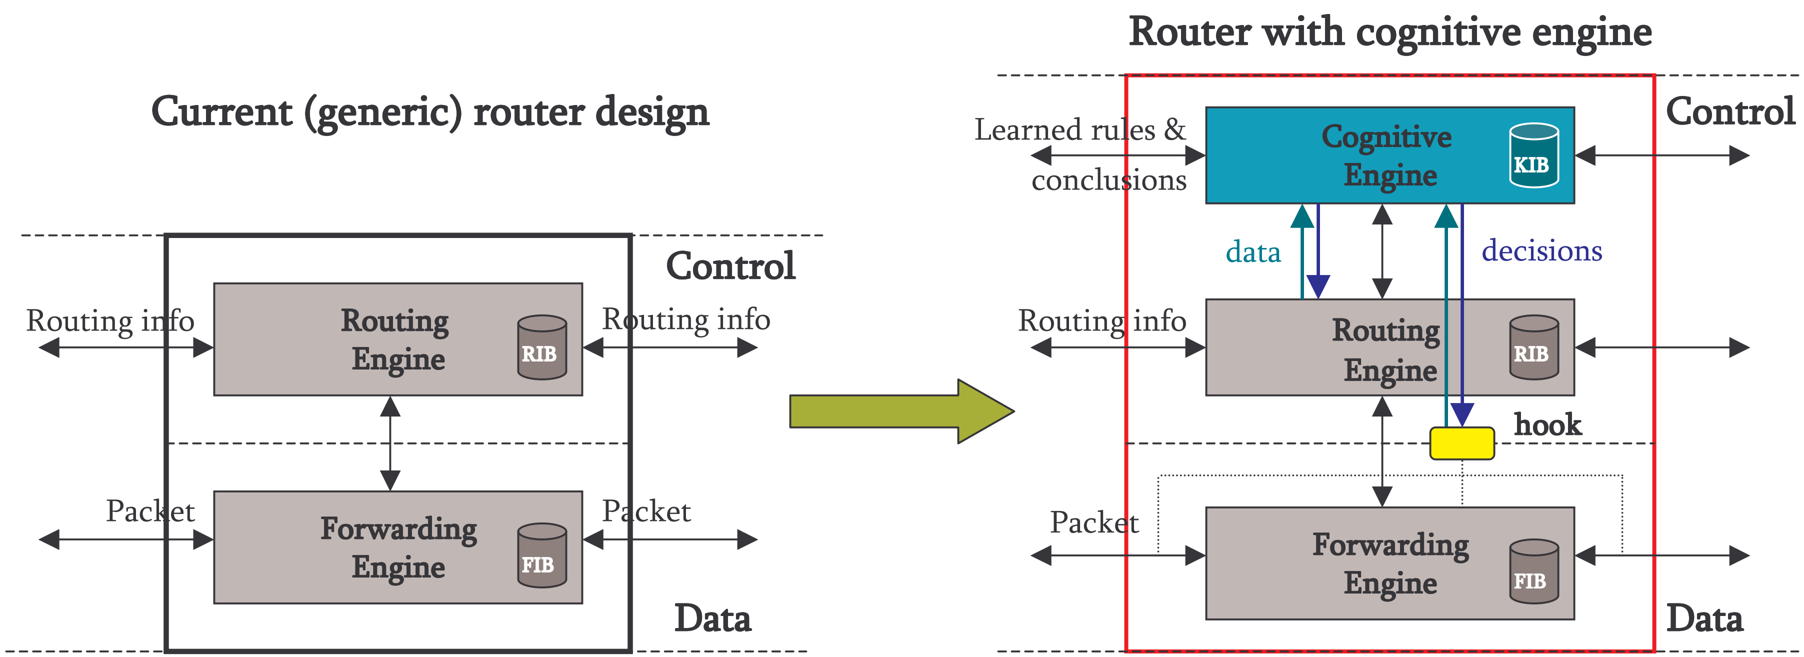

Objective:

Augment existing routing/control paradigm of (system & network) lower-level data collection and decision making, with a cognitive engine that

- Enables system & network to learn about its own behavior and environment over time
- Analyzes problems, tunes its operation and increases its functionality and performance

Cognitive engine using semi-supervised, online, and distributed machine learning

Example: User-network cooperation

Connectivity selection among multiple and multi-homed sites

Example: <s,d_1> via [{1,2,3} → {4,5,6,7}] or <s,d_2> via [{1,2,3} → {7,8}]

-> **select (s1,d1) over (s2,d2)**

Edge routers path performance monitoring (passive or active) -> extract information from monitoring data (note: multiple cycles) so as to provide path quality prediction

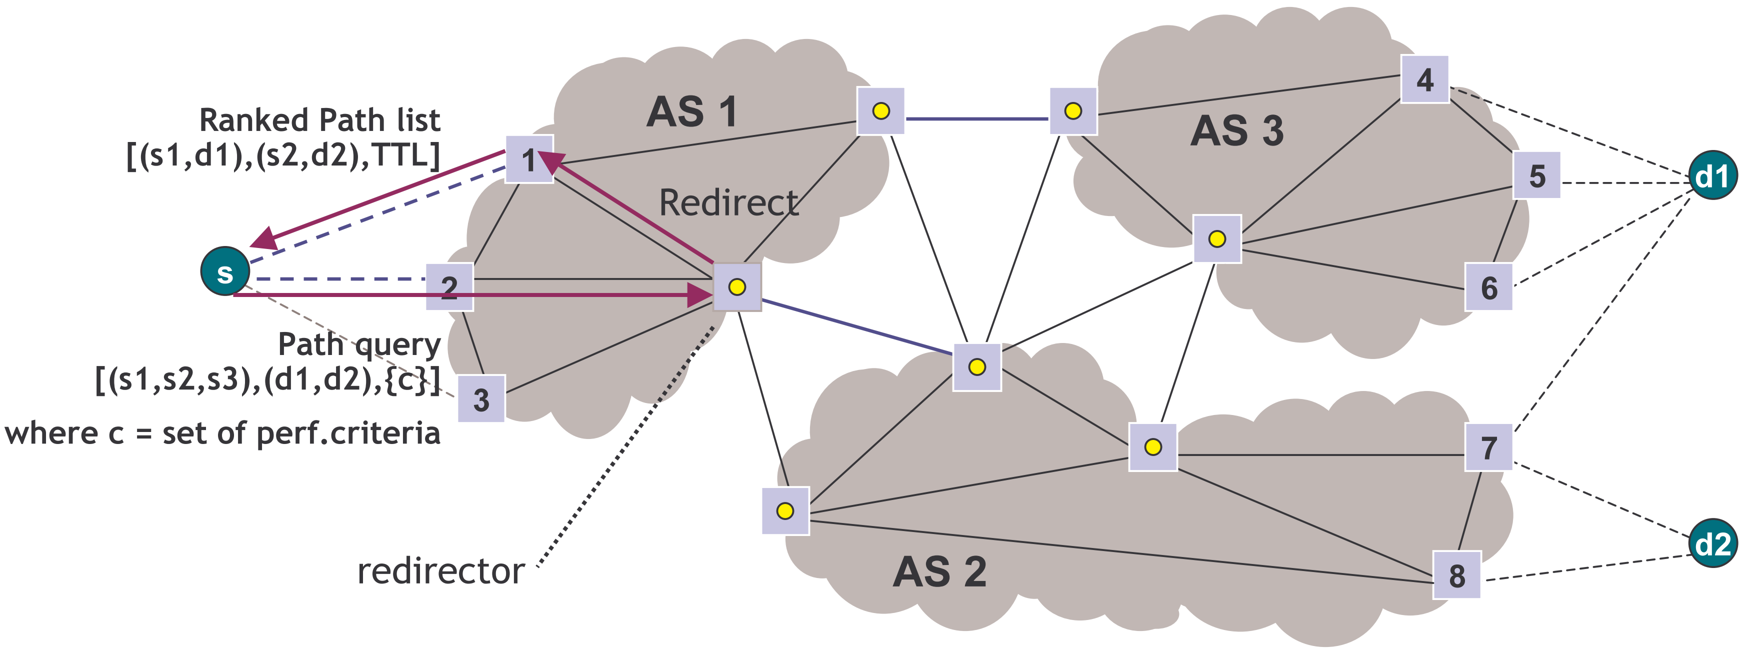

Knowledge Plane (KP) - [Clark'03]

- Driving idea: abstract/isolate high-level goals from low-level actions

  Augment control system with a higher-level structure that addresses issues of "knowing what is going on" in the network
- Unified KP
  - Common standards/framework for "knowledge"
  - Structured based on knowledge, not task
- Objective: build a new network generation
  - Drive its own deployment and automate its (re-)configuration by learning from past to improve future performance
  - Diagnose its own problems (with imperfect and conflicting info)
  - Make defensible decisions about how to solve them (respond to problems/attacks in better timeframe than manual intervention)
  - Recognize/mediate conflicts in policies and goals

Knowledge Plane (KP) - Architecture

- **Core foundation**: ability to integrate behavioral models and reasoning processes into a networked environment
- Separate structure (breaking boundaries of legacy control system) sitting on top of the current control system
- Comprising cognitive tools and learning
  - Sensors: produce observations
  - KP: Process + Data structures: PIB - K-base - Models
  - Actuators/effectors: to change/alter network behaviour / environment

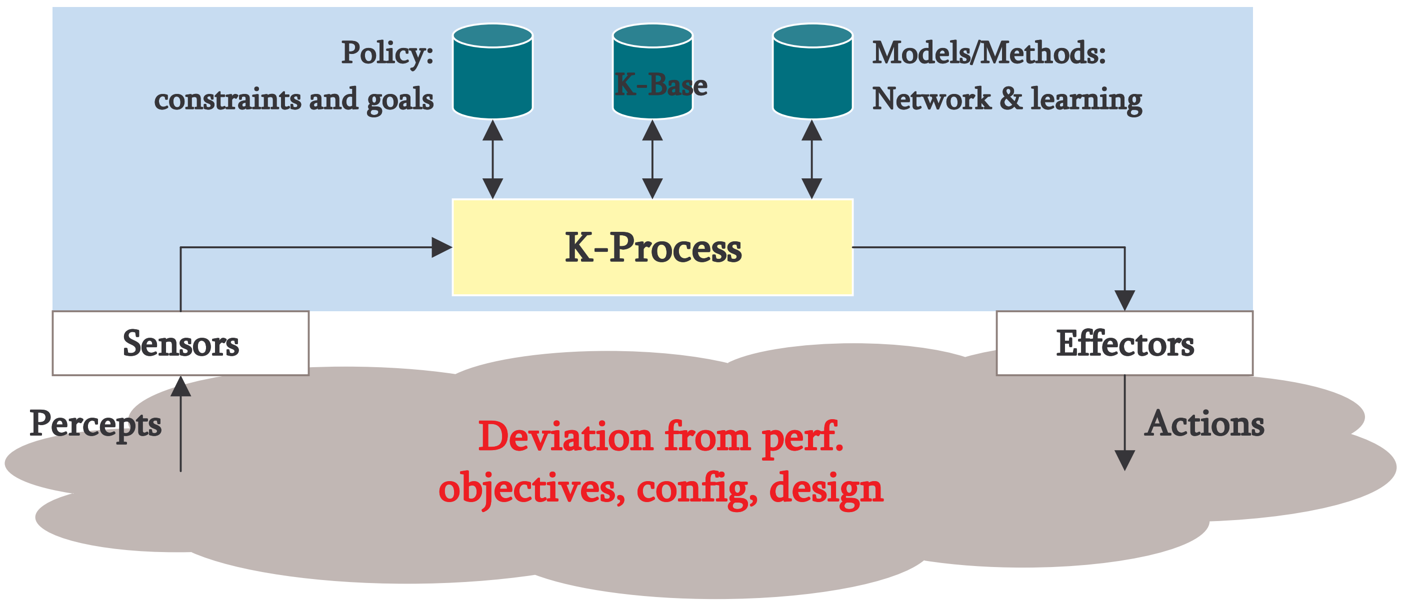

Positioning against KP

- Cognitive routing system should
  - Be modular instead of relying on a unified approach for development and deployment reasons
    - Examples: access vs core, edge vs intermediate router
  - Rely on network relative view (i.e. forwarding and routing) rather than network global view

    Reason: prevent scaling issues, increase resiliency, and organic deployability

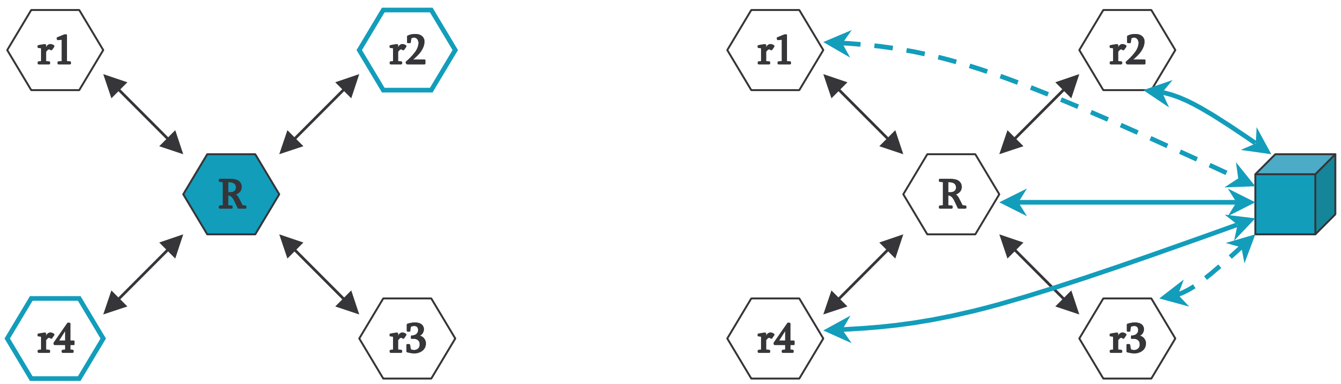

Positioning against KP

- Cognitive routing system should
  - Be architected taking into account inherent distributed properties and capabilities of forwarding and routing system
    - Example: intra- vs inter-domain routing
    - Example: flow vs aggregate forwarding

    Instead of being constructed as a uniform and ubiquitous two-dimensional structure (plane)

    Reason: developability (sys.engineering) and adaptability

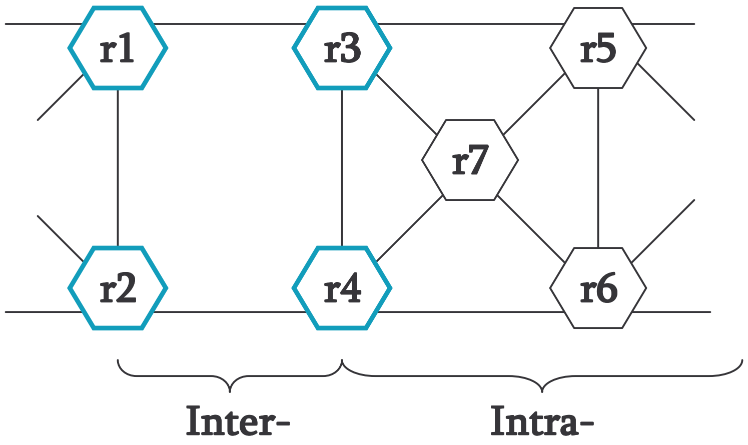

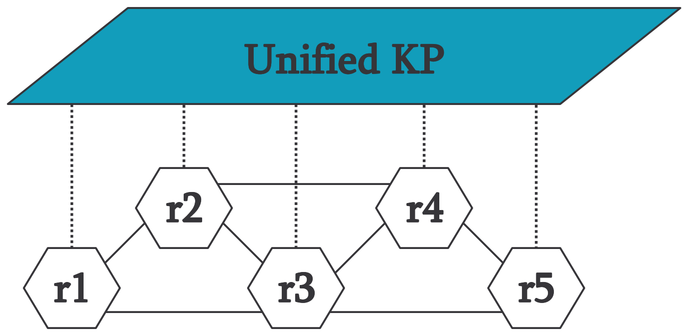

Positioning

- Cognitive routing system should
  - semi-supervised instead of supervised learning techniques
  - on-line instead of off-line learning techniques
  - distributed instead of centralized learning techniques

Challenges & Perspectives

- Set of networking use cases representative of future Internet challenges
- Applying machine-learning methods (using cognitive engine designed as set of functional blocks) to these use cases
- Experimental evaluation (using physical and virtual facility)

Technical Challenges	Use Case
Adaptive traffic sampling and management, path performance monitoring, and intrusion and attack/anomaly detection	Adaptive traffic sampling and management
	Path performance monitoring
	Cooperative intrusion and attack/anomaly detection
Path availability, network recovery and resiliency, and profile-based accountability	Path availability
	Network recovery and resiliency
	Profile-based accountability
Routing system scalability and quality	Routing system scalability and routing system quality (convergence, stability/ robustness, and stretch)

→ Cognitive engine building blocks (architecture) and low-level components

Challenges & Perspectives

- Cognitive system
  - Components, levels of interaction & coupling between components
  - Communication protocols
    - Internal: between internal components
    - External: between cognitive engines
- Performance gain
  - Semi-supervised, and online learning
  - Semi-supervised, online, **AND** distributed learning

    Determine predictive value (->decision), and derive appropriate set of commands (directed to routing and forwarding engine)
- Collected data translation & representation as information to cognitive engine
  - Determine impact on lower-level cognitive engine components

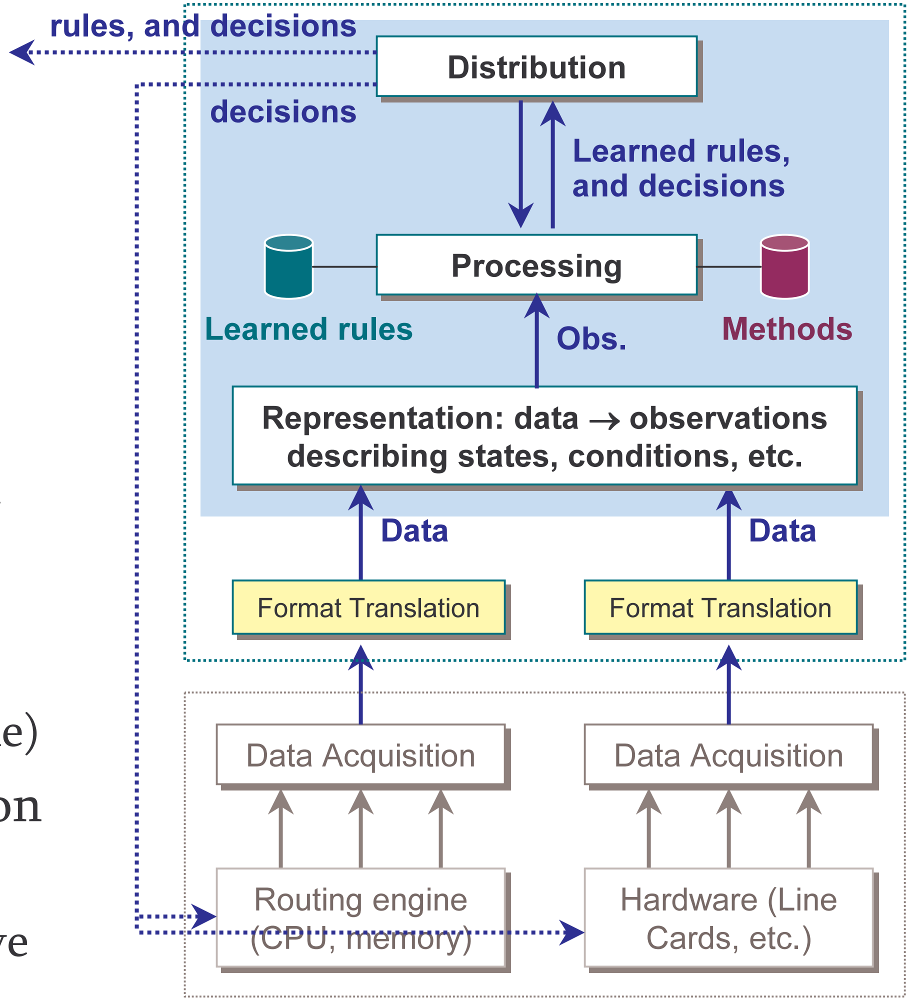

Challenges & Perspectives

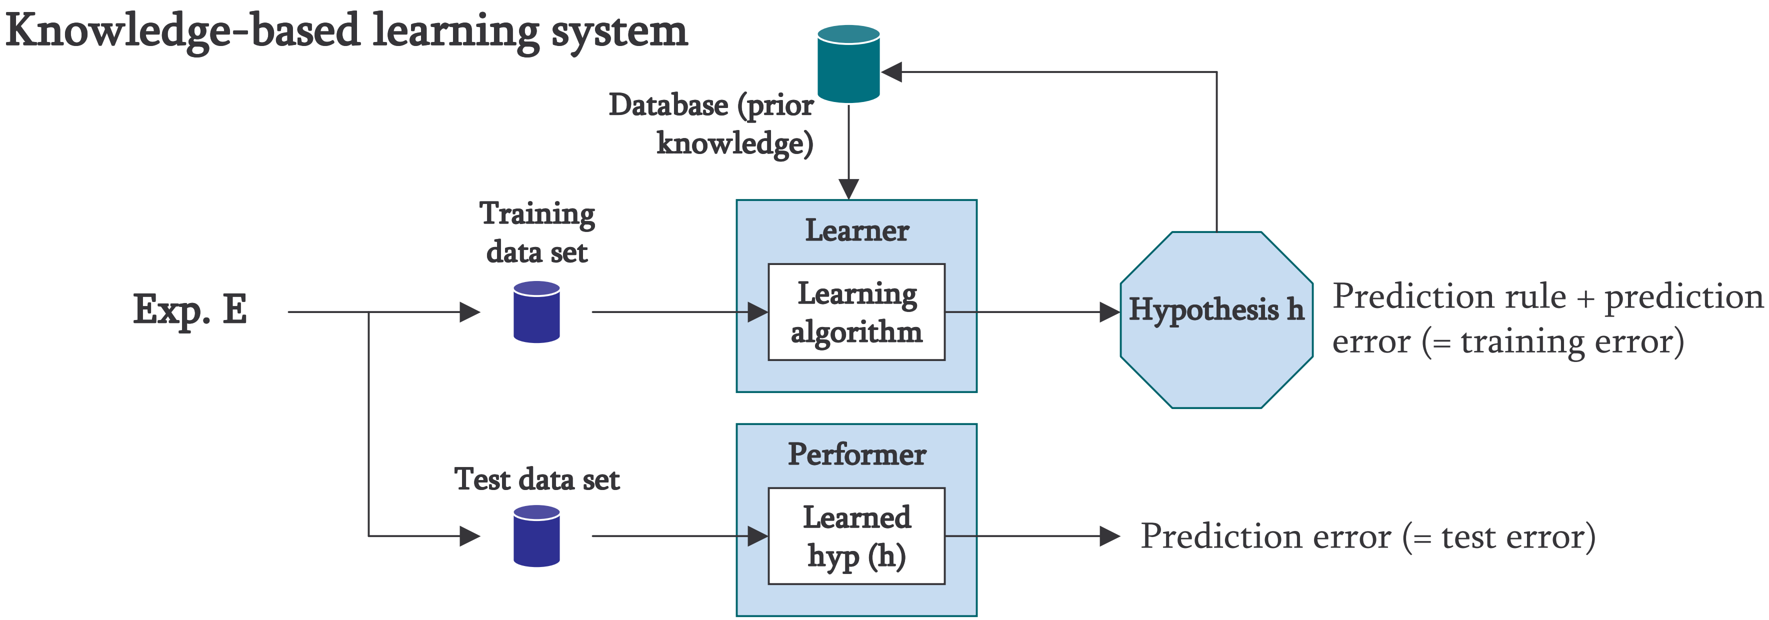

- **Algorithmic**
  - Which learning algorithms (on-line & distributed) ?
  - How well do these algorithms perform?
- **Training data/examples**
  - How much training data is sufficient to learn a task/model with high confidence?
  - Are some training examples more useful than others?
- **Knowledge**
  - When is it useful to use prior knowledge?
  - What is the best way to represent and utilize knowledge?
  - How to distribute/disseminate (route) knowledge?

Conclusions

Introduction of a new architectural component

-> cognitive routing system:

1. Improve and extend the Internet functionality by providing adequate solutions to the existing and foreseeable upcoming Internet challenges
2. Limit infrastructure and operational cost and complexity resulting from Internet growth (compared approach(es) consisting in continuously patching existing routing equipment)
3. Ensure Internet viability by removing complexity, from existing components, but adding functionality