

Simulation Semantics

A computational framework linking language,
cognition, and action

Srini Narayanan

ICSI and UC Berkeley

snarayan@icsi.Berkeley.edu

Staff Research Scientist

Google Zurich

srinin@google.com

Language, Learning and Neural Modeling

www.icsi.berkeley.edu/AI

- Scientific Goal
 - Understand how people learn and use language
- Practical Goal
 - Deploy systems that analyze and produce language
- Approach
 - Build models that perform cognitive tasks, respecting all experimental and experiential constraints
 - Cognitive linguistic theories with advanced biologically-based computational methods



Real language understanding

- Brazil's disappointing economy stuck in the mud and on the verge of sinking deeper. No more room for the Government to maneuver. (The Economist, June 8, 2013)



Language Learning

..Of the above possible fields the learning of languages would be the most impressive, since it is the most human of these activities. This field seems however to depend rather too much on sense organs and locomotion to be feasible.

– Alan Turing, Intelligent Machinery, Sept. 1947

Conceptual Metaphor

- Many abstract concepts have conventional metaphorical conceptualizations
 - normal everyday ways of using concrete concepts to reason systematically about abstract concepts.

Ex. Event Structure Metaphor

- Most abstract reasoning makes use of embodied reasoning via *metaphorical mappings from concrete to abstract domains*



Simulation hypothesis

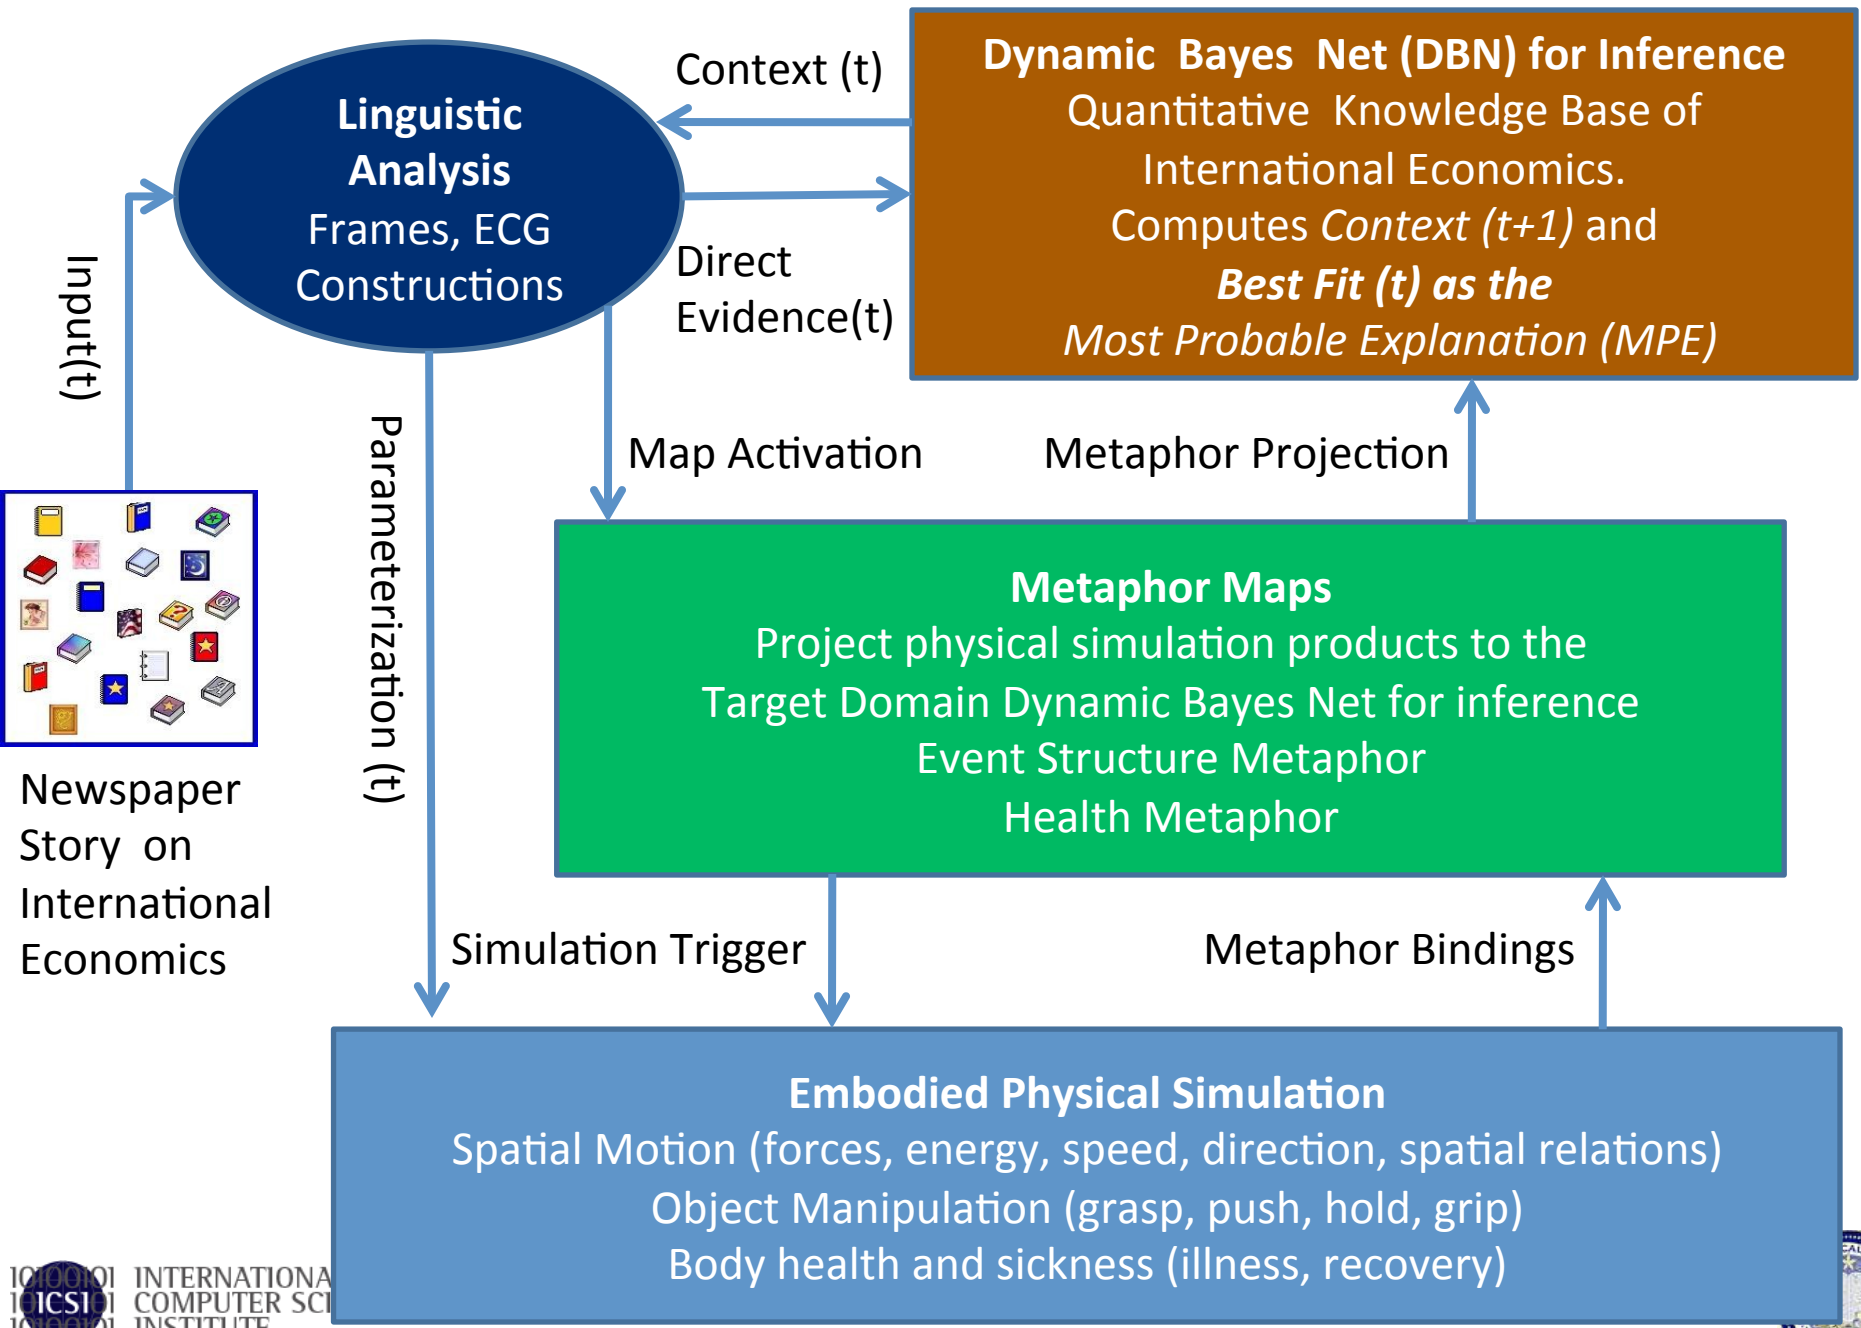
(N1997, FN2004, F2006, LN2010, in prep)

We understand utterances by mentally simulating their content.

- Simulation exploits some of the **same neural structures** activated during performance, perception, imagining, memory...
- Linguistic structure **parameterizes** the simulation.
 - Language in context gives us enough information to simulate
 - Linguistic structure includes constructions, frames (<http://framenet.icsi.berkeley.edu>), embodied schemas, metaphor, mental spaces. Formalized in ECG.
 - Understanding involves running the *best fitting* simulation for the linguistic input.
 - Analysis involves finding the best simulation specification (Bryant 2008).

Pilot Task: Interpret simple discourse fragments/ blurbs

- US Economy **on the verge of falling back** into recession after **moving forward** on an **anaemic recovery**.
- Indian Government **stumbling** in implementing Liberalization plan.
- **Moving forward** on all fronts, we are going to be **ongoing** and **relentless** as we **tighten the net** of justice.
- The Government **is taking bold new steps**. We are **loosening the stranglehold** on business, **slashing** tariffs and **removing obstacles** to international trade.



Uses of simulation

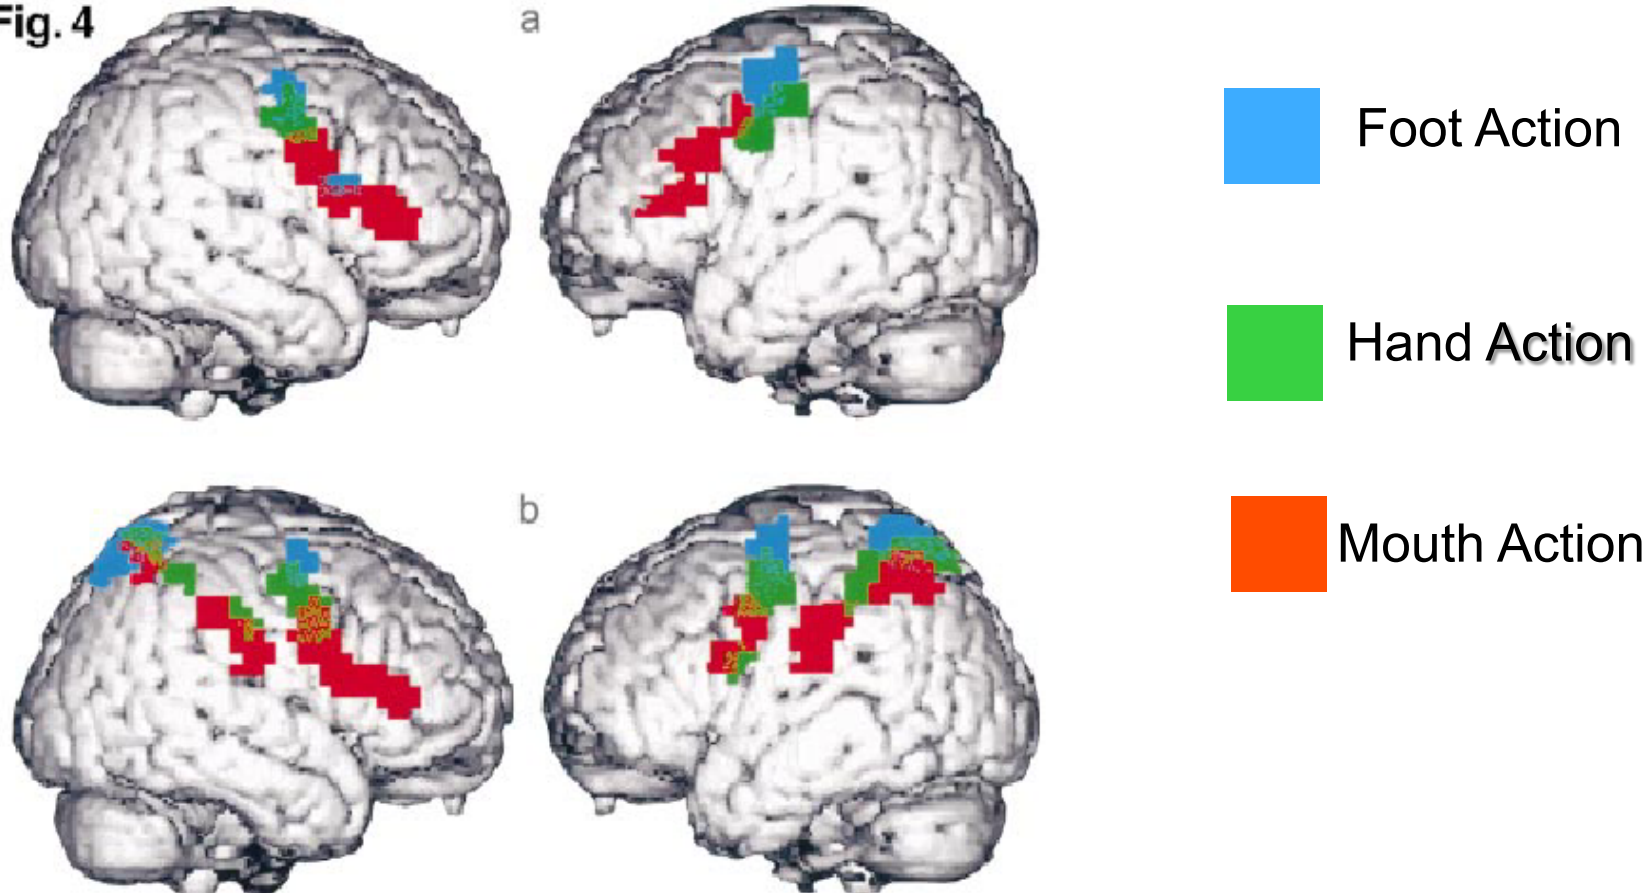
- Results show that motion related constructions often the most effective method to provide the following types of information about abstract plans and actions.
 - Information about **uncertain events and dynamic changes in goals and resources**. (sluggish, fall, off-track, no steam)
 - Information about **evaluations of policies** and economic actors and **communicative intent** (strangle-hold, bleed).
 - Communicating **complex, context-sensitive and dynamic** economic scenarios (stumble, slide, slippery slope).
 - Communicating complex **event structure and aspectual information** (on the verge of, sidestep, giant leap, small steps, ready, set out, back on track).

Simulation appears necessary for tractable inference about events and processes.

(Narayanan 99, Sinha and Narayanan 2005, Gedigian et al., 2006, Sinha 2008, Lakoff and Narayanan 2010, Narayanan 2012)

Simulation and mirror systems

Fig. 4



Buccino et al., (2002), European journal of neuroscience

Aziz-Zadeh et al. (2007), Social Neuroscience

Bergen et al, (2008), Cognitive Science

Aziz-Zadeh et al. (2013), Society for Neuroscience

Active/Procedural Semantics circuits versus program semantics

AMBLING



TROTting



GALLOPING

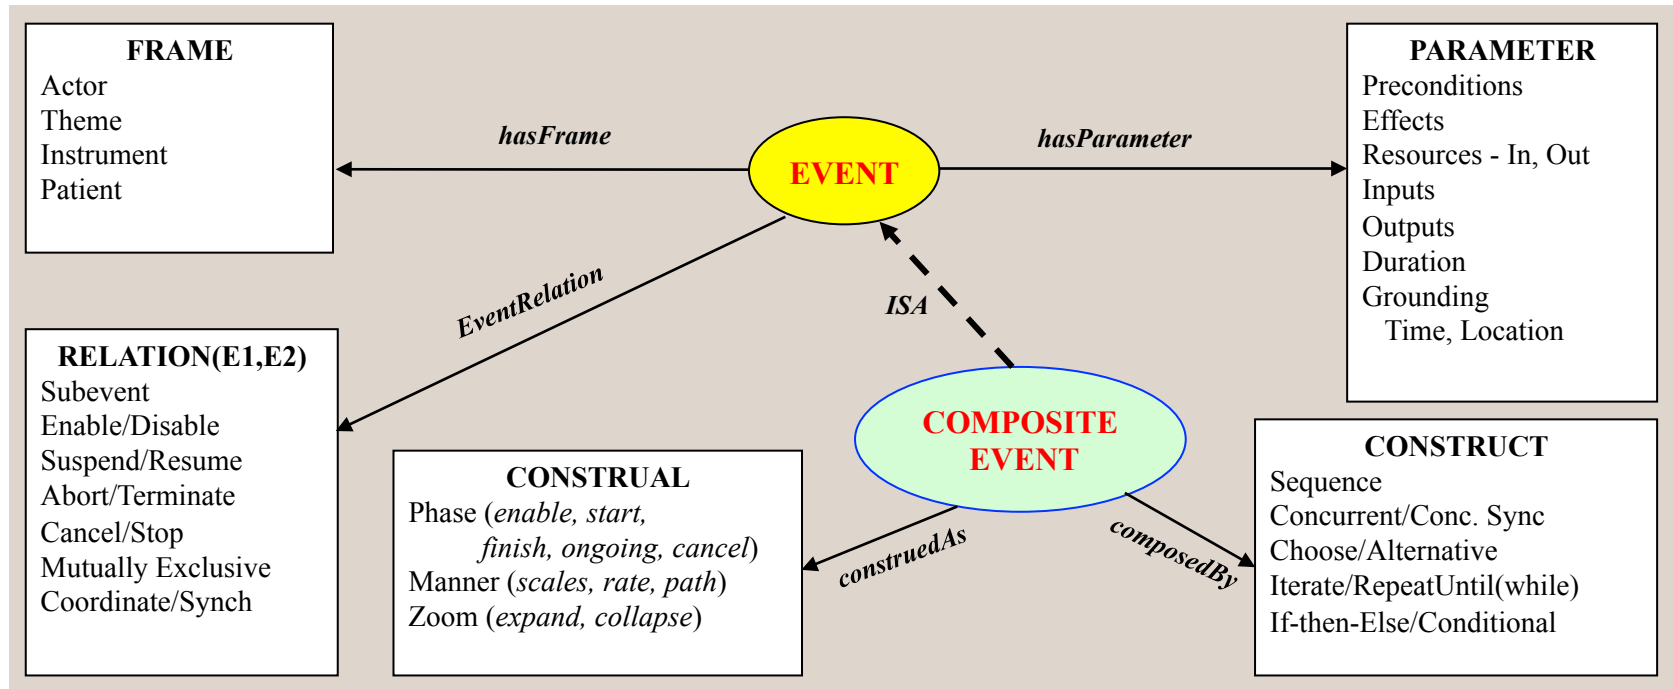


ROBERTO C&T

- Actions are:
 - Parameterized
 - Interruptable
 - Coordinated
 - Hierarchical
 - Relational
 - Have hybrid dynamics
 - Multiple scales
 - Discrete/continuous

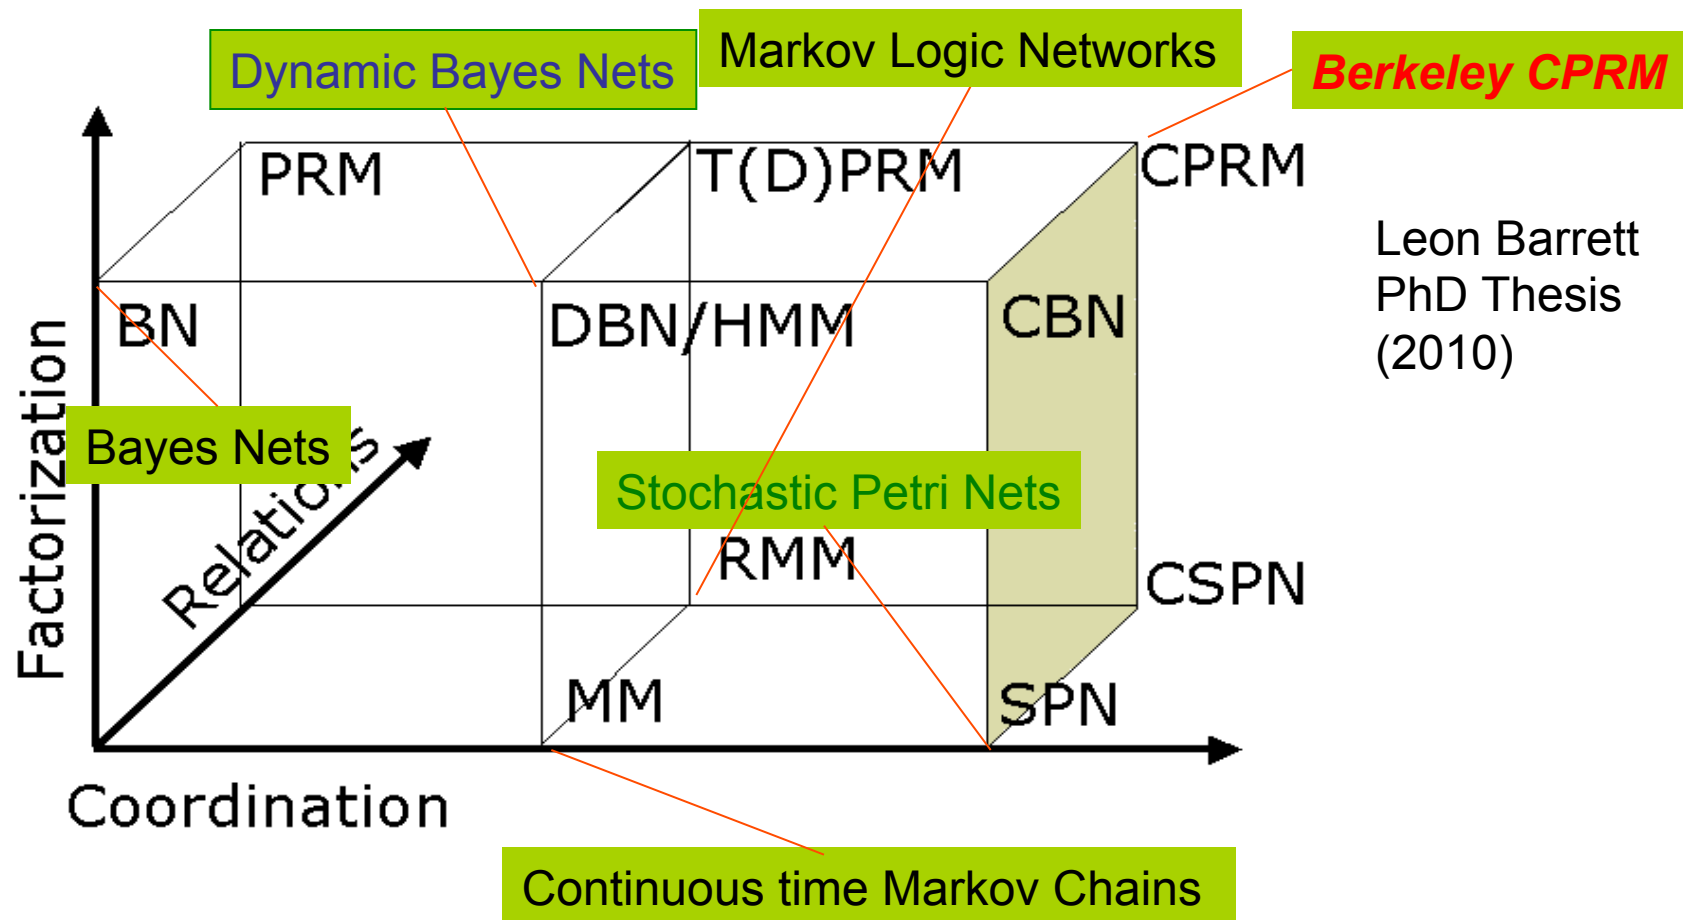
- Procedural semantics can be code
 - Inference with and learning code are both intractable
- *XNETS: Graphical Model for action and simulation.*

The structure of Events: A Formalized event schema



- Key elements
 - preconditions, resources, effects, sub-events
 - evoked by frames (alternatively: predicates, words)
- Contrast with Event Recognition/Extraction, other NLP work
 - [Bethard '07], [Chambers '07,09][McCloskey 11]

Structured Probabilistic Inference with Events



Leon Barrett
PhD Thesis
(2010)

Probabilistic inference

– Filtering

- $P(X_t | o_{1..t}, X_{1..t})$
- Update the state based on the observation sequence and state set

– MAP Estimation

- $\text{Argmax}_{h_1..h_n} P(X_t | o_{1..t}, X_{1..t})$
- Return the best assignment of values to the hypothesis variables given the observation and states

– Smoothing

- $P(X_{t-k} | o_{1..t}, X_{1..t})$
- modify assumptions about previous states, given observation sequence and state set

– Projection/Prediction/Reachability

- $P(X_{t+k} | o_{1..t}, X_{1..t})$

Current State: The Language of Events

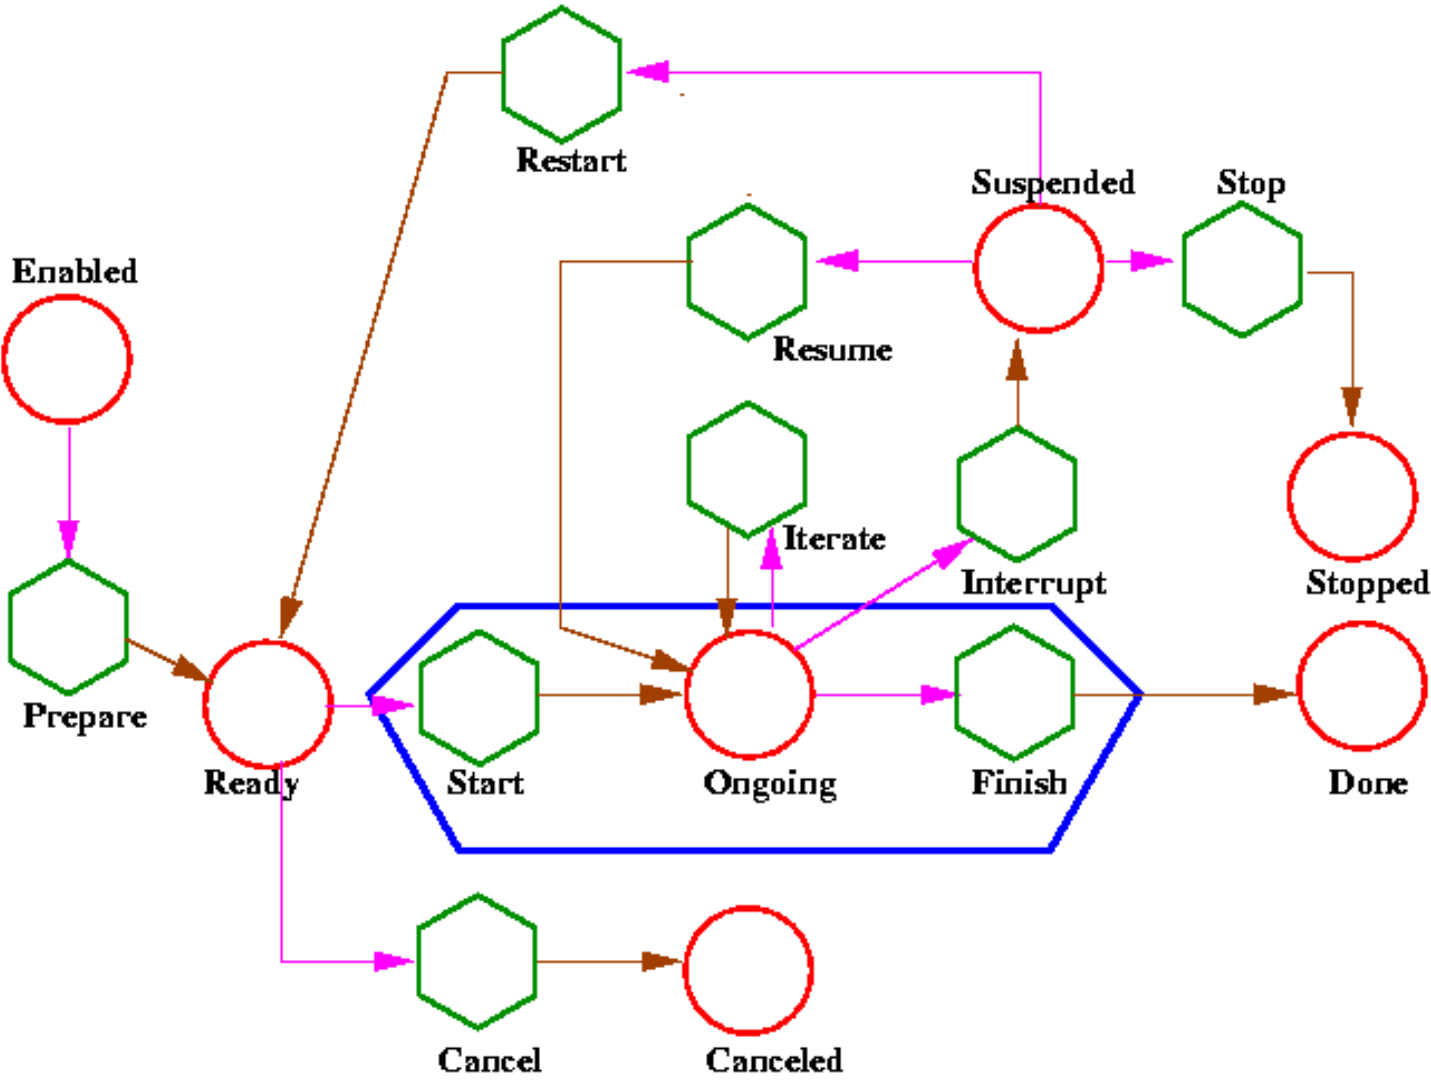
- Event Frames
- *Aspect/Tense*
- Affordances
- Modals
- Causation/Force Dynamics
- Counterfactuals
- Scales
- *Event Structure Metaphor*

Cognitive Science 2013, ICLC 2013, Neural Computation 2013, Cambridge Encyc. Of Psycholinguistics 2012, AAAI SS 2013, CompBio 2013, NAACL 2013(S), *sem2013 (S), CL 2012 (S)

Aspect and Event Structure

- Aspect is the name given to the ways languages describe the **structure of events** using a variety of **lexical and grammatical** devices.
 - **Viewpoints**
 - is walking, walk
 - **Phases of events**
 - Starting to walk, walking, finish walking
 - **Inherent Aspect**
 - run vs cough vs. rub
 - **Composition with**
 - Temporal modifiers, tense..
 - Noun Phrases (count vs. mass) etc..

The Controller X-Net And Event Structure



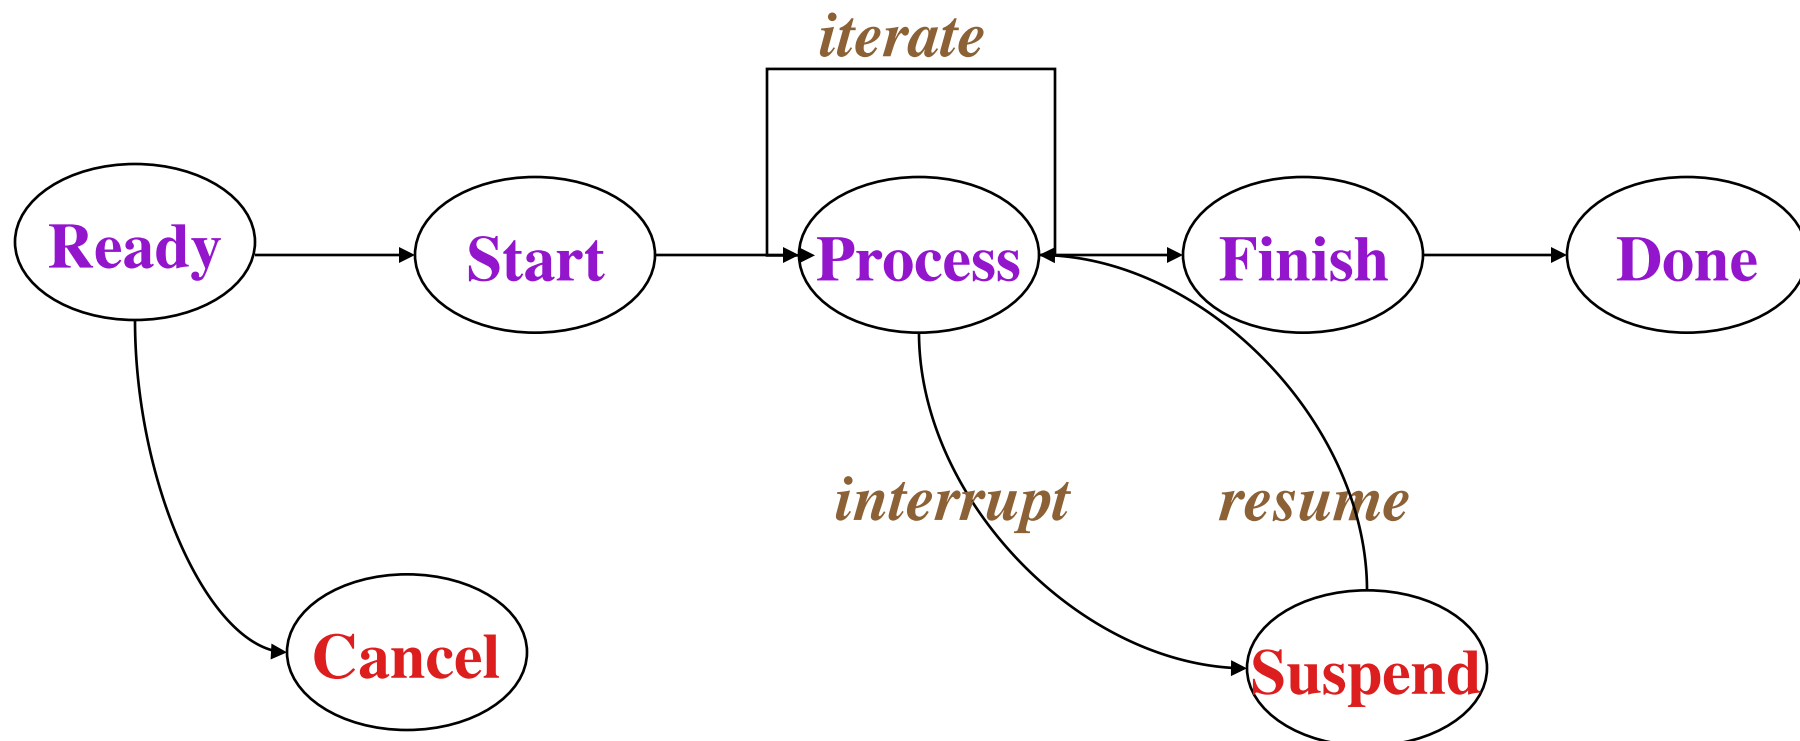
Phases, Viewpoints, and Aspects

- John **is walking** to the store.
- John **is about to walk** to the store.
- John **walked** to the store.
- John **started walking** to the store.
- John **is starting** to walk to the store.
- John **has walked** to the store.
- John **has started to walk** to the store.
- John **is about to start** walking to the store.
- John **resumed walking** to the store.
- John **has been walking** to the store.
- John **has finished walking** to the store.
- John **almost** walked to the store.

Aspectual Information in Questions

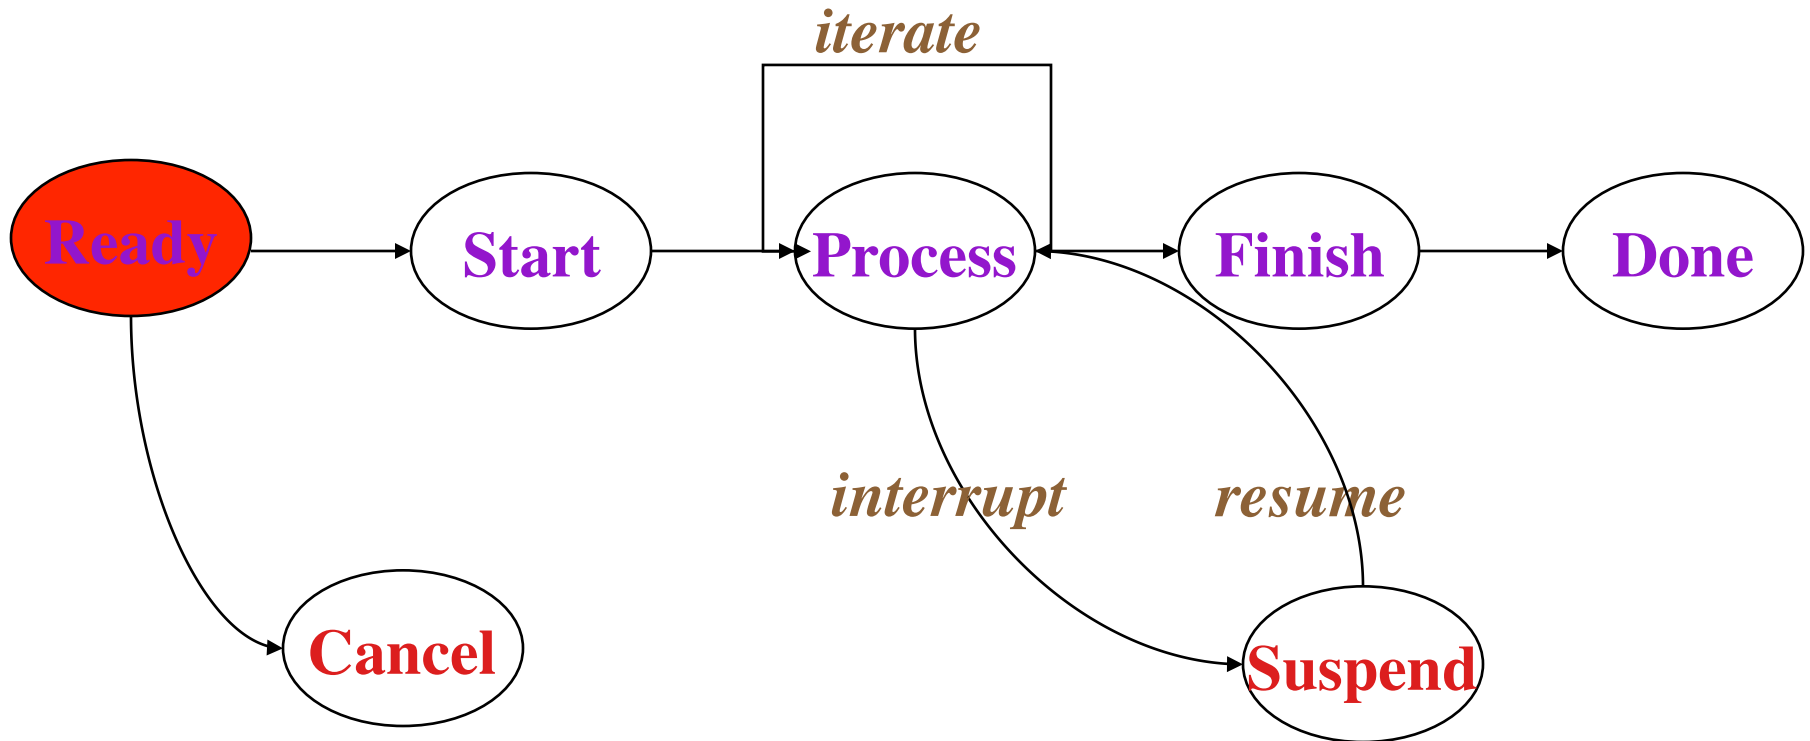
- The US **continues to be concerned** about Russian arms sales to Syria. Has the US **taken any steps to try to stop** the sale of Russian arms to Syria?
- The US **is concerned** about Russian arms sales to Syria. Has the US **taken any steps to try to detect** the sale of Russian arms to Syria?
- The US **is starting to be concerned** about Russian arms. Is the US **ready to take steps** to try to establish the sale of Russian arms to Syria?

A Generic Process Schema



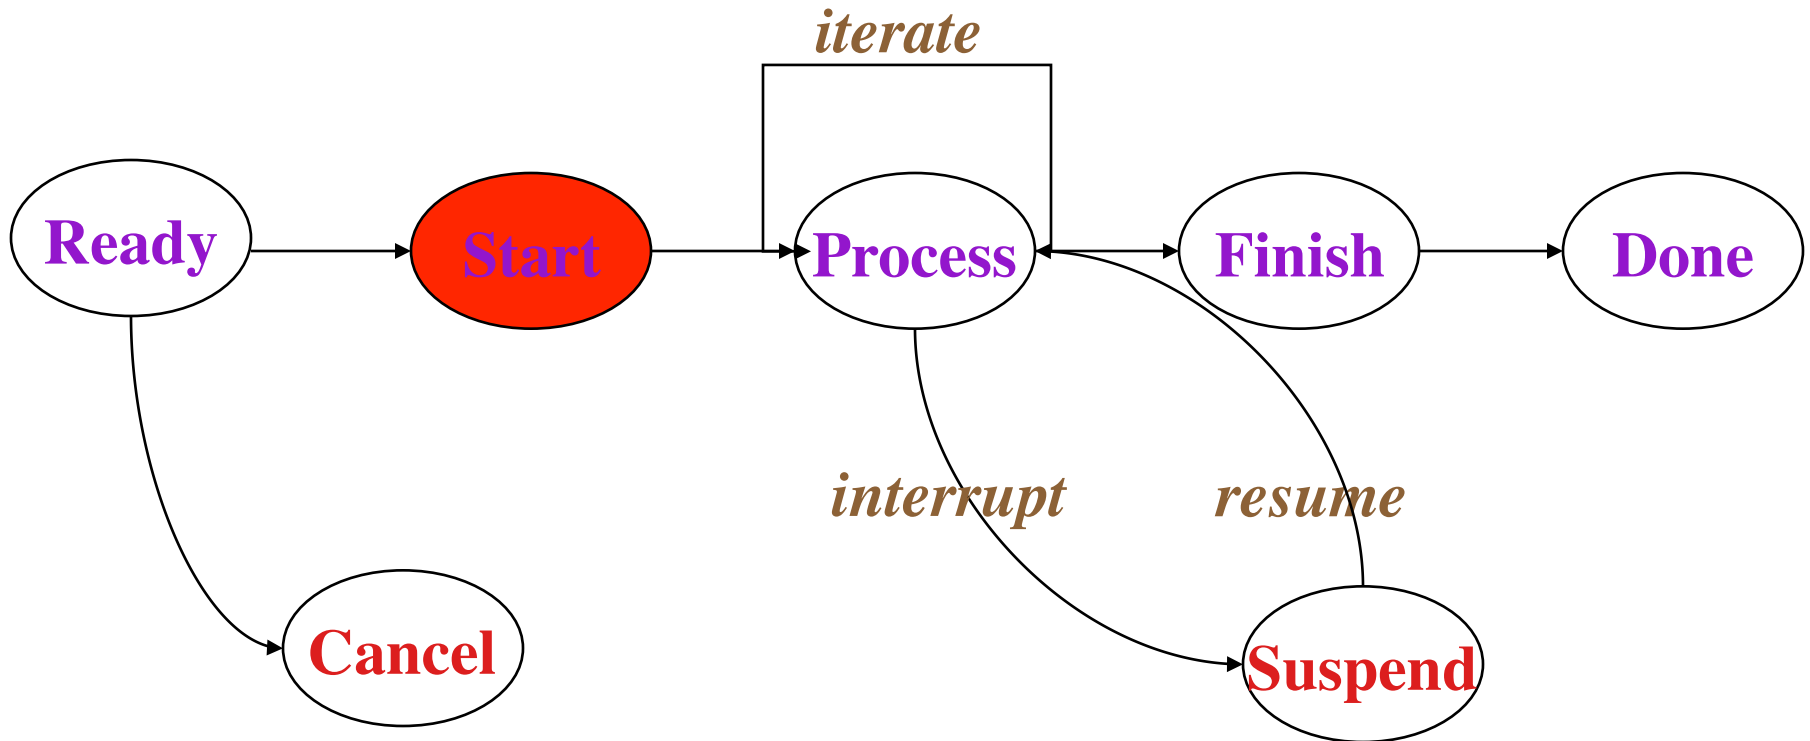
- Part of **Conceptual Structure**.
- Generalizes over actions and events. Has **internal state** and **models evolution of processes**.

About to X



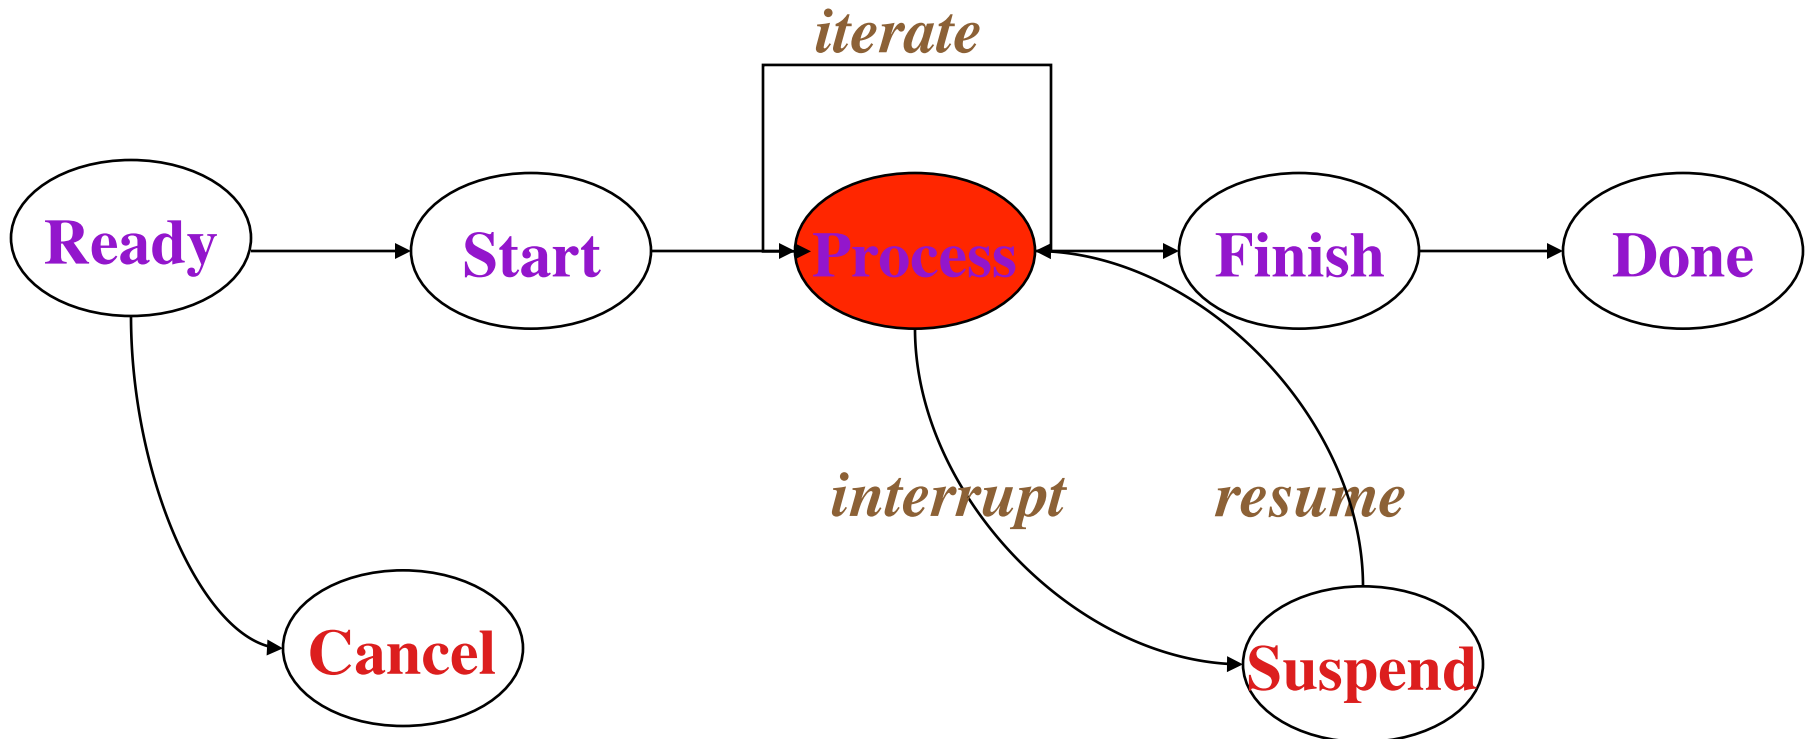
- **Part of Conceptual Structure.**
- **Generalizes over actions and events. Has internal state and models evolution of processes.**

Starting to X



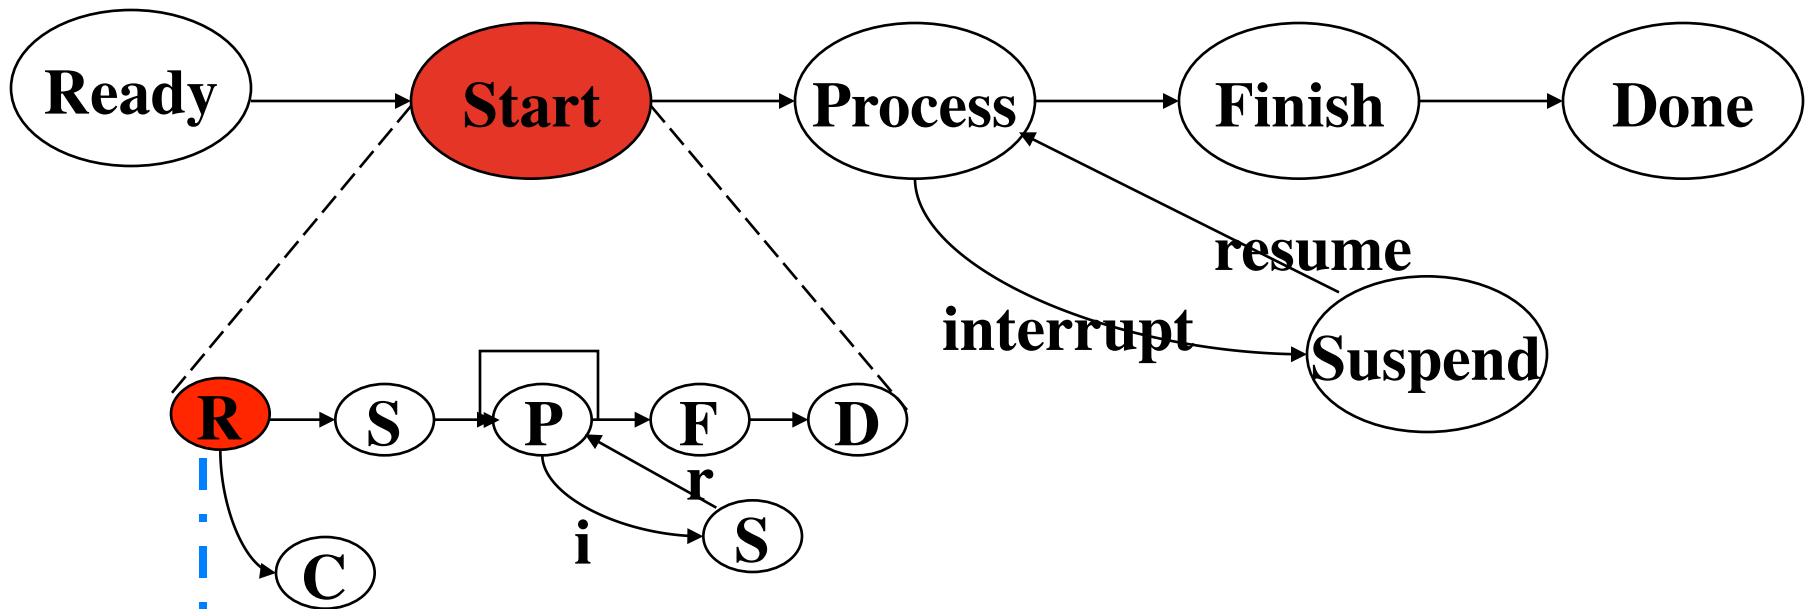
- **Part of Conceptual Structure.**
- **Generalizes over actions and events. Has internal state and models evolution of processes.**

is X-ing



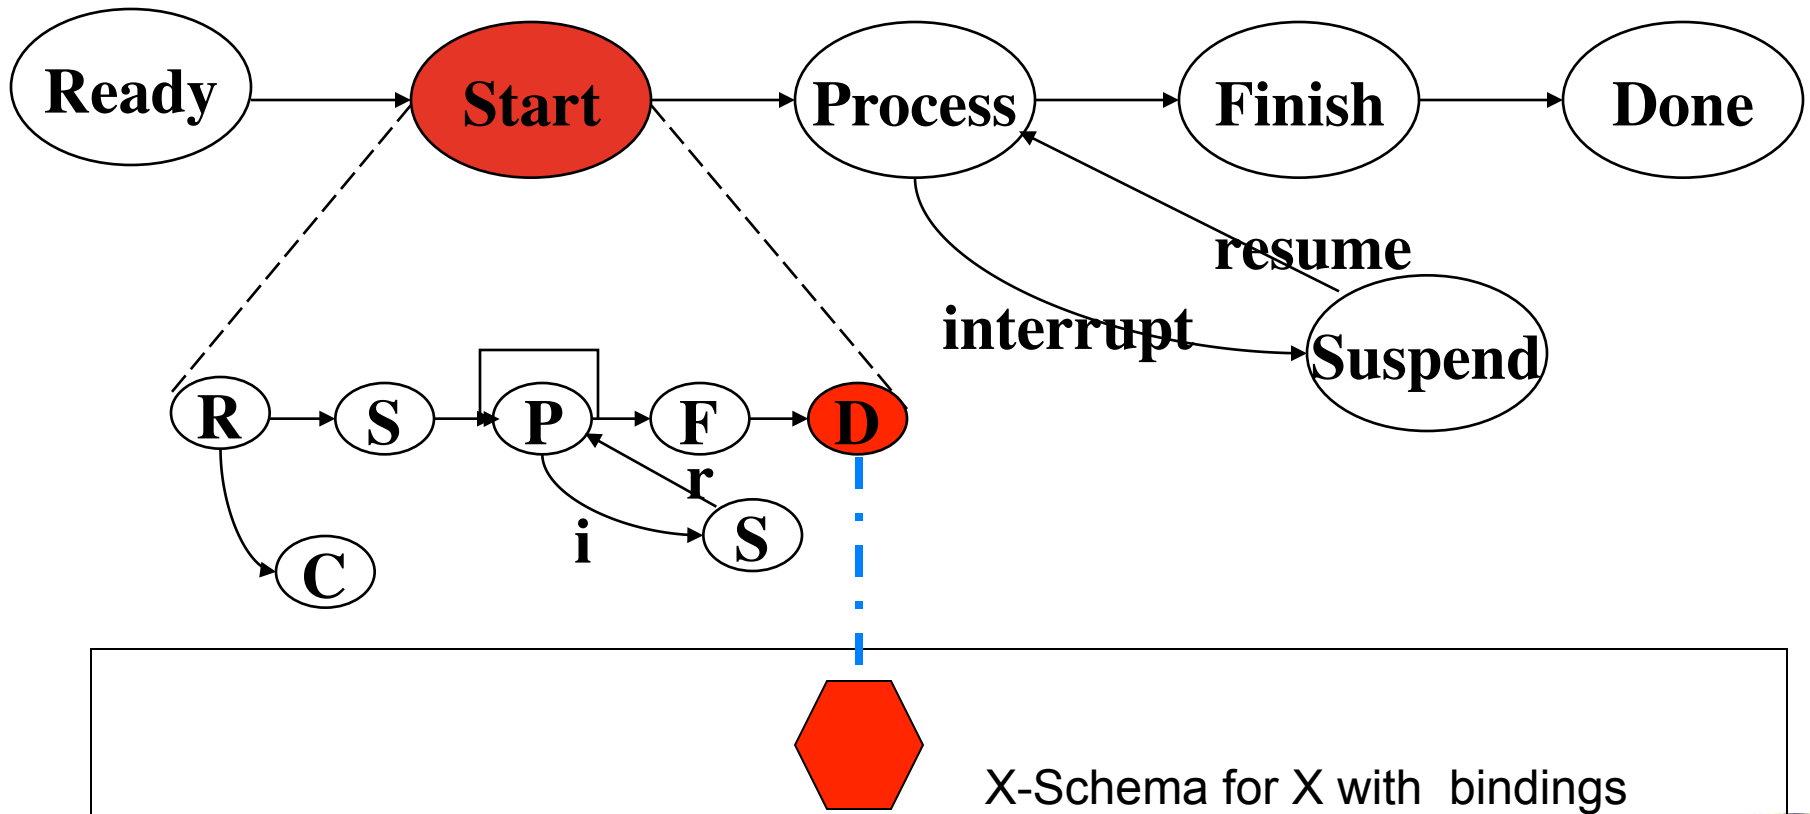
- Part of **Conceptual Structure**.
- Generalizes over actions and events. Has **internal state** and **models evolution of processes**.

Embedding: About to start (X)



X-Schema for X with bindings

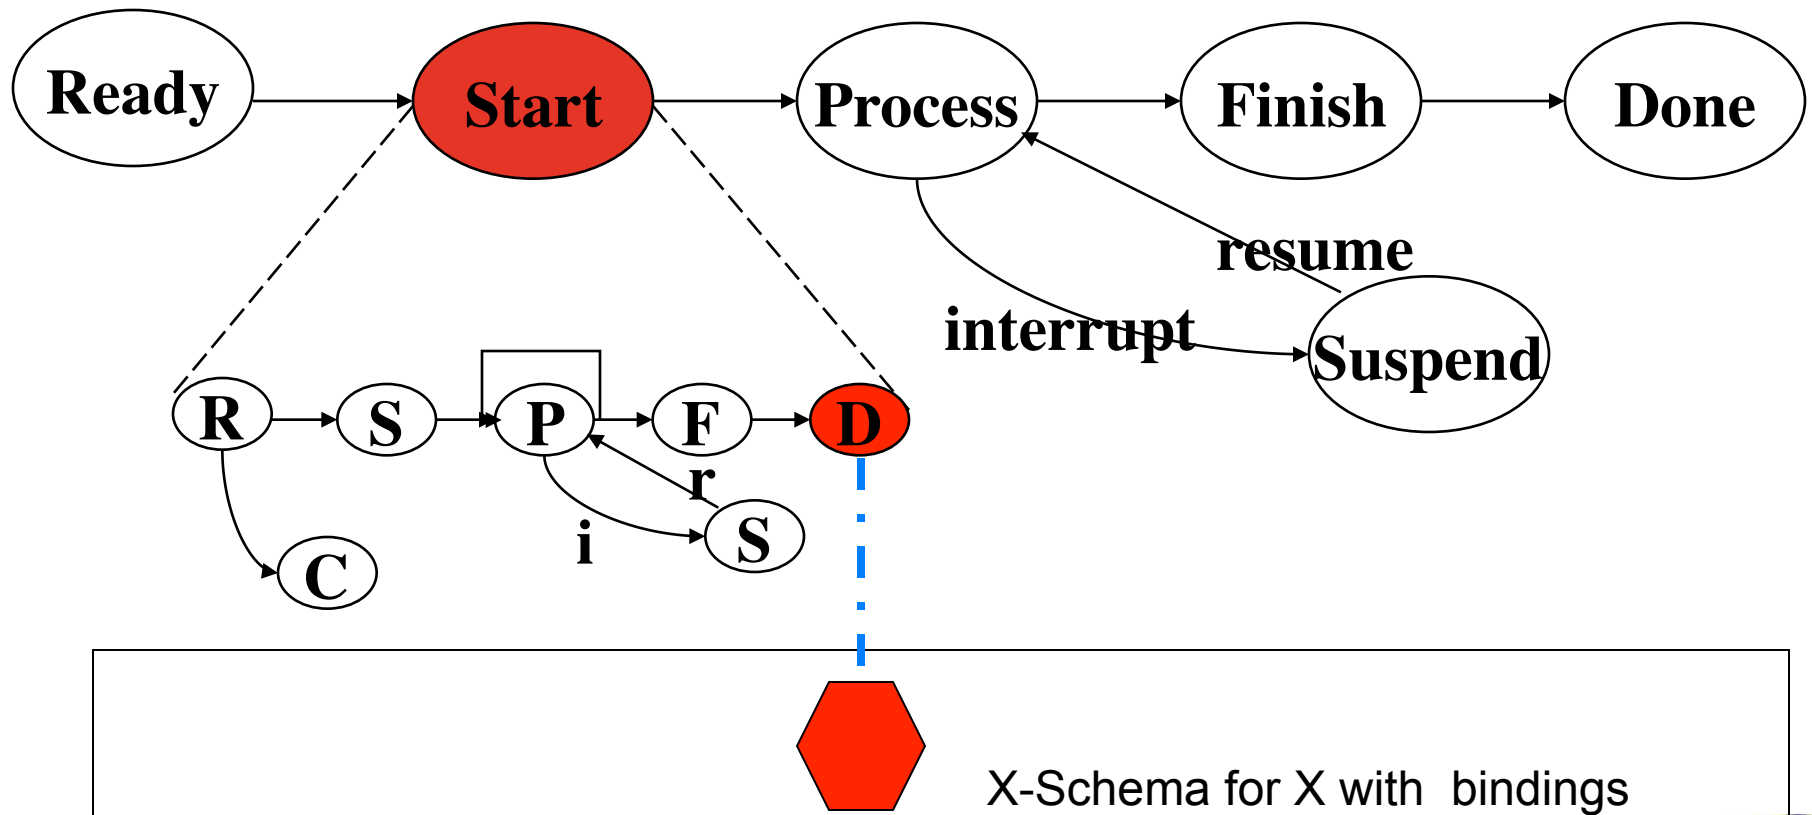
Embedding: Has Started (to X)



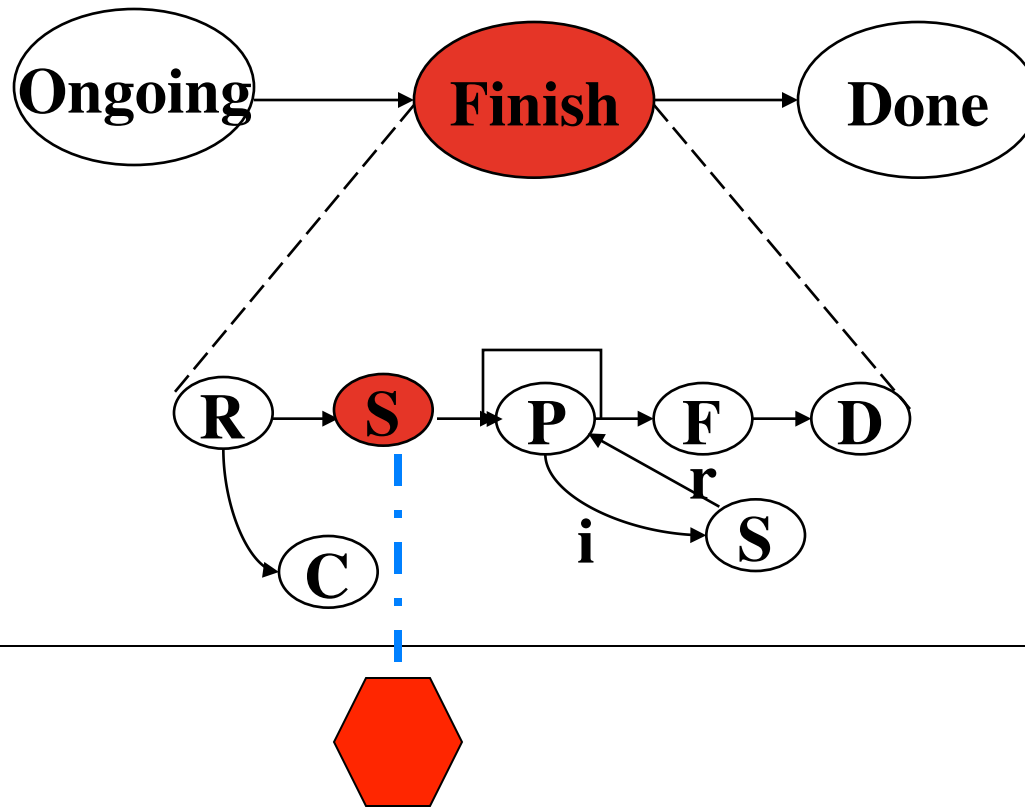
What does this mean?

- Now this is not the end. It is not even the beginning of the end. But it is, perhaps, the end of the beginning.
 - [Winston Churchill](#)

Embedding: The end of the beginning



Embedding: The beginning of the end



X-Schema for X with bindings

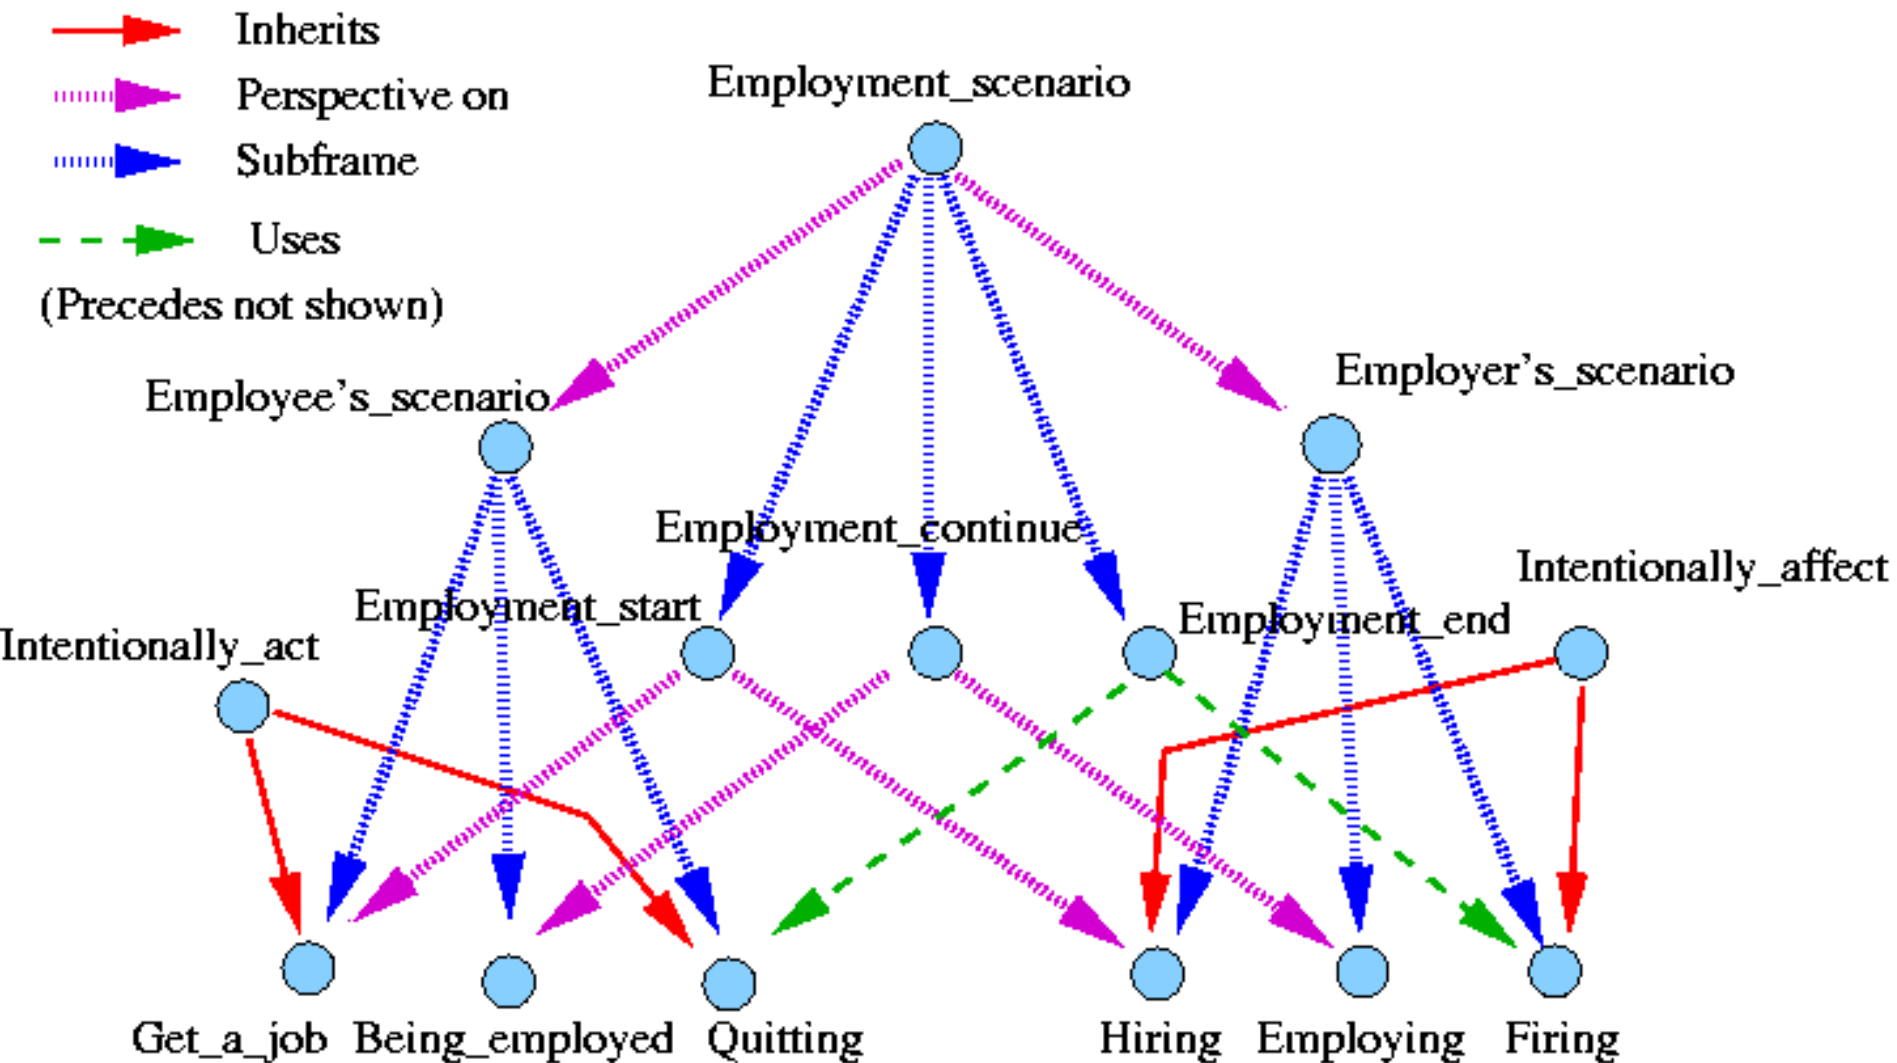
General and Domain Knowledge

- Conceptual Knowledge and Inference
 - Embodied
 - Language and Domain Independent
 - Powerful General Inferences
 - Ubiquitous in Language
- Domain Specific Frames and Ontologies
 - *FrameNet*, ontologies, entities, lexicon, lexical-relations
- Metaphor links domain specific to general
 - E.g., France slipped into recession.
 - *MetaNet*

The FrameNet Project

- FrameNet is a lexical resource organized around Semantic frames: events, relations, and states which are the basis for understanding groups of word senses, e.g. the Being_employed frame contains *work.v*, *position.n*, *employed.a*, *jobless.a*, etc.
- Frames are distinguished by the set of roles involved, known as **frame elements**, in this case, **Employee**, **Employer**, **Field**, **Place of employment**, etc.
- Sentences are annotated to exemplify these FEs, e.g.
[Employee She] [Time recently] *accepted* [Contract_basis part-time]
WORK [Employer at ICSI].
- FN currently contains > 1,110 frames and 170,000 annotations

The FrameNet Project



MetaNet

Goal: to build a system that extracts metaphors from text in four different languages

English, Persian, Spanish, Russian

Purpose: To understand the role metaphor plays in how people from different cultural backgrounds make judgments and decisions



Combining Multiple Systems

- Dual systems
 - Conceptual Semantics:
 - Construction based system for LM detection
 - Mapping to Sources through Schemas and Frames
 - Affect identification with source and target frames
 - Distributional Semantics
 - Seed based semi-supervised system for LM detection
 - Mapping to sources through subcategorization and distributional information

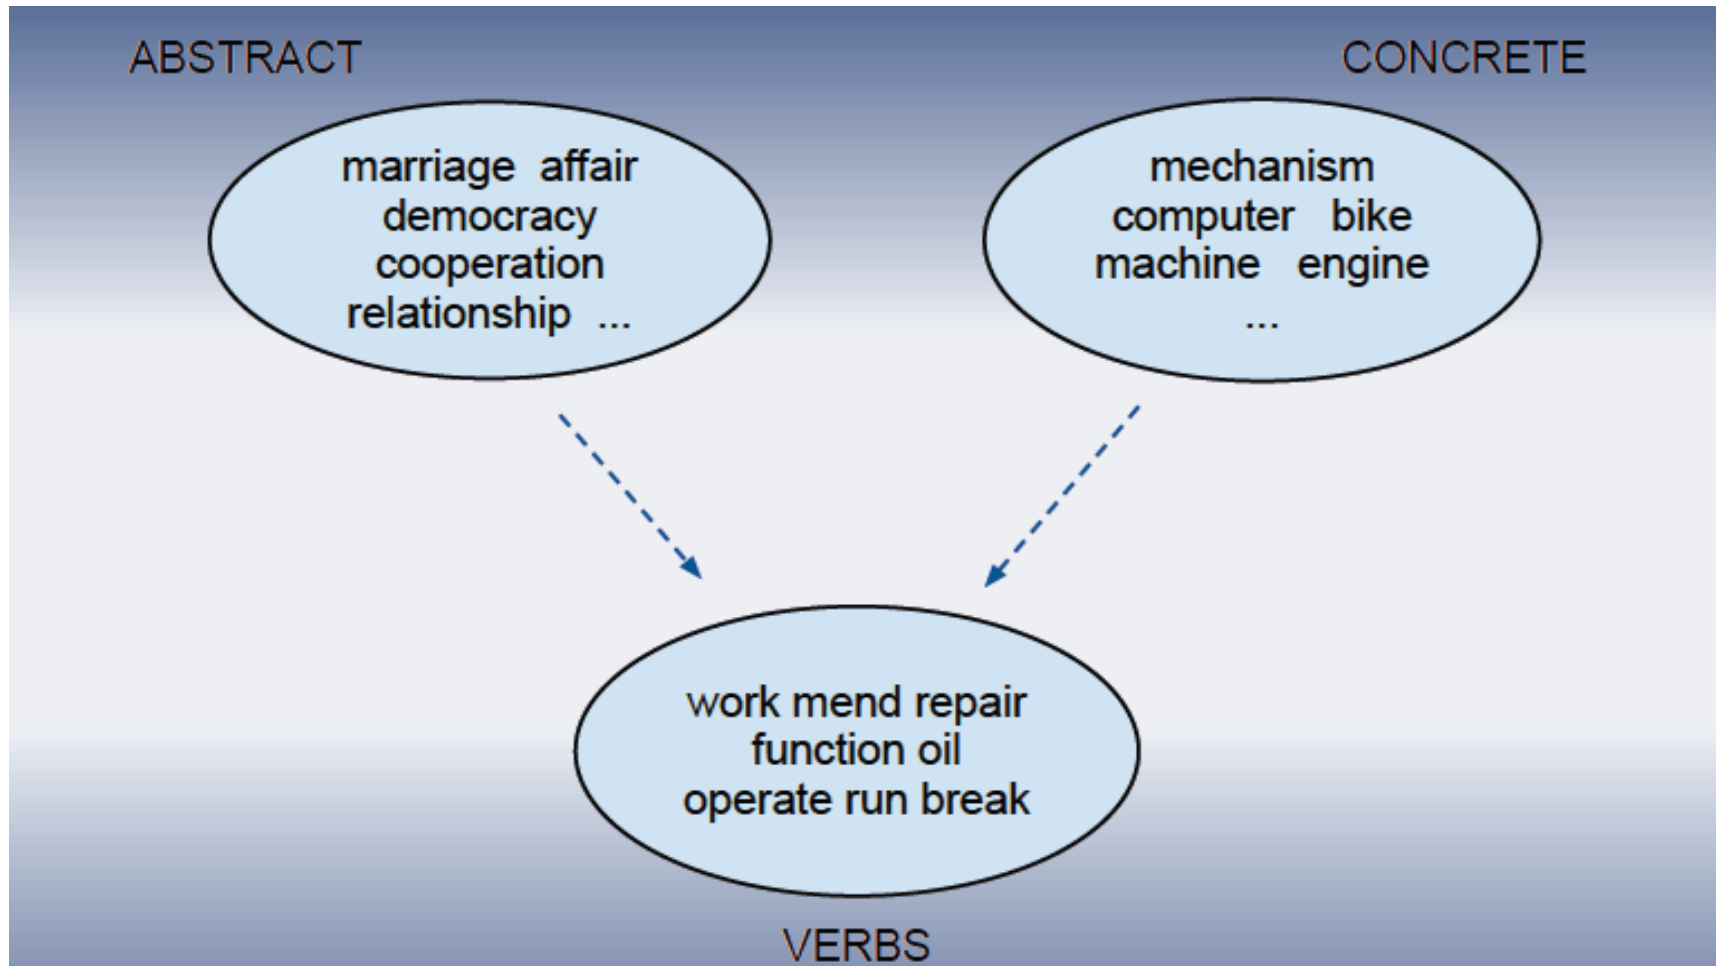
Evaluations

- Detection of LMs
- Mapping LMs to CM Target and Source Concepts
 - Source dimensions
- Hybrid systems
 - Detection:
 - Construction Matching System (CMS)
 - Language Model System (LMS, Persian)
 - Mapping/source:
 - Semantic Cluster Overlap System (SCOS)
 - Distributional mapping system
- Computing of affect

Experiments and Systems

- Semi-supervised learning (spectral clustering) to discover linguistic metaphors (LMs)
- Unsupervised learning (HGFC) to mine conceptual metaphors from large amounts of linguistic data
- Selectional preferences and a vector space model to assign the corresponding conceptual metaphors (CMs) and source dimensions to LMs
- Multilingual topic modeling experiments to detect cross-cultural differences
- Linguistic annotation experiments and supervised learning to model affect of LMs.

Spectral Clustering over dependency relation fillers (Shutova)



Examples

Seed phrase expansion

face deficit → tackle poverty
combat poverty → cure unemployment
expose poverty → disguise inequality
destroy resource → attack capital etc.

Sentences containing the identified metaphors

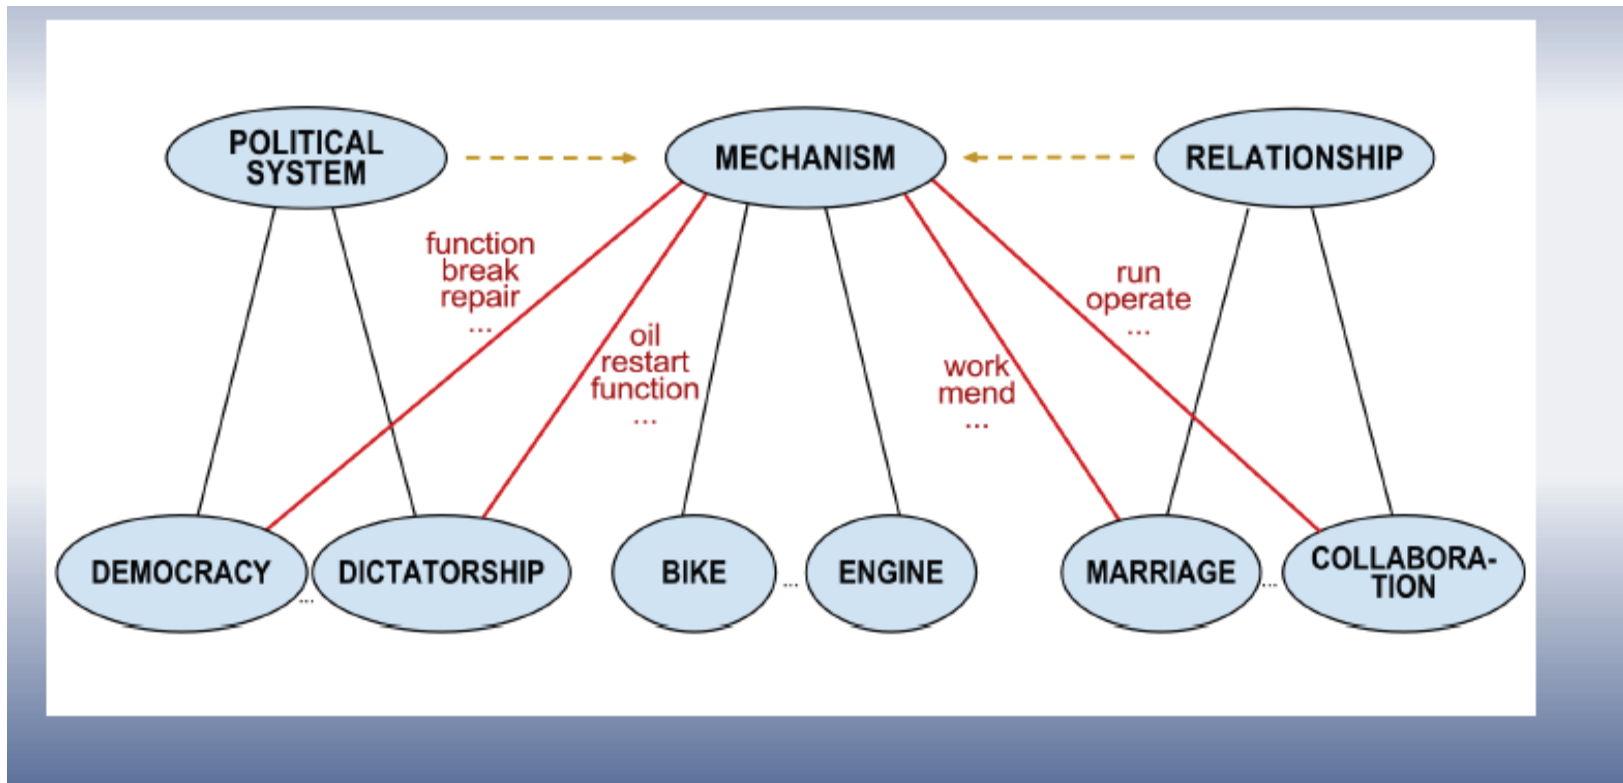
Tackling poverty will require a multi-faceted approach and a coordinated effort by social service agencies, non-profit organizations, and community volunteers.

[..] ways to **cure unemployment** after college graduation.

For a time, asset bubbles **disguised inequality**, [..]

Precision 0.79

Hierarchical Factorized Graph Clustering



Example Output

SOURCE: fire

TARGET: sense hatred emotion passion enthusiasm sentiment hope interest feeling resentment optimism hostility excitement anger

TARGET: coup violence fight resistance clash rebellion battle drive fighting riot revolt war confrontation volcano row revolution struggle

SOURCE: disease

TARGET: fraud outbreak offence connection leak count crime violation abuse conspiracy corruption terrorism suicide

TARGET: opponent critic rival

FEELING IS FIRE LMs

anger *blazed* (Subj), optimism *raged* (Subj), enthusiasm *engulfed* them (Subj), passion *flared* (Subj), interest *lit* (Subj), *fuel* resentment (Dobj), anger *crackled* (Subj), feelings *roared* (Subj), *light* with hope (lobj)

CRIME IS A DISEASE LMs

cure crime (Dobj), abuse *transmitted* (Subj), *eradicate* terrorism (Dobj), *suffer from* corruption (lobj), *diagnose* abuse (Dobj), *combat* fraud (Dobj), *cope with* crime

LM Detection Performance

- Tuned to favor precision
 - EN Recall=0.836 (153/183) Precision=0.793 (153/193)
 - ES Recall=0.804 (115/143) Precision=0.891 (115/129)
 - FA Recall=0.484 (44/91) Precision=0.623 (44/104)
 - RU Recall=0.545 (67/123) Precision=0.971 (67/69)
- Tuned to favor recall
 - EN Recall=0.847 (155/183) Precision=0.718 (155/216)
 - ES Recall=0.846 (121/143) Precision=0.691 (121/175)
 - FA Recall=0.527 (48/91) Precision=0.558 (48/134)
 - RU Recall=0.626 (77/123) Precision=0.846 (77/91)

LM/CM Mapping Performance

- m4mapping
 - EN Recall=0.760 (139/183) Precision=0.777 (139/179)
 - ES Recall=0.748 (107/143) Precision=0.775 (107/138)
 - FA Recall=0.407 (37/91) Precision=0.536 (37/69)
 - RU Recall=0.675 (83/123) Precision=0.748 (83/111)
- m4source
 - EN Recall=0.710 (130/183) Precision=0.710 (130/183)
 - ES Recall=0.622 (89/143) Precision=0.627 (89/142)
 - FA Recall=0.308 (28/91) Precision=0.406 (28/69)
 - RU Recall=0.504 (62/123) Precision=0.525 (62/118)

Experimental Research

- Role of moral systems in grounding metaphors of governance and economic inequality
- Up-Down scalar effects of verticality metaphors of economic inequality
- Decision making with metaphors of emotion
- Motion (manner, aspect) and economic inequality metaphors
- Metaphor choice and inference

State of the Repository

- Implemented for all four languages:
 - RDF triplestore repository, Semantic MediaWiki environment, LM extraction, SQL export (DEMO)
 - Automatically Extracted Mappings
 - 1200 Vetted Conceptual Mappings (Frame to Frame)
 - 10000 Linguistic metaphors (lexicalized mappings)
 - ~150000 example annotations
- Connection to inference
 - Probabilistic Network Analysis
 - Inference through simulation
- Initial results in several ICLC 2013 papers, NAACL 2013, ACL Workshop 2013, 2014.

Metaphor Visualizer

Active Ontology | Entities | Classes | Object Properties | Data Properties | Individuals | DL Query | Metaphor Visualizer

Metaphor Visualizer View Component:

Currently Selected Metaphor:

CORRUPTION IS A DESTRUCTIVE NATURAL PROCESS

Metaphor Selection

Metaphors | Schemas

CORRUPTION.EFFECT IS DISEASE.EFFECT
CORRUPTION IS A DESTRUCTIVE NATURAL PROCESS
 CORRUPTION IS A DISEASE
 CORRUPTION IS A TERMITE INFESTATION
 CORRUPTION IS DECAY
 CORRUPTION IS EROSION
 CORRUPTION.SCALE IS DISEASE.SCALE
 COURSE OF A LIFE IS THE PATH TRACED BY A L
 CREATING A LAW IS GROWING A PLANT
 CREATING IS MAKING
 CRIME IS AN ANIMAL
 CRIME IS A DISEASE

Relationships Selection (1 or more)

Metaphors | Schemas

isRelatedToMetaphor
 entailsMetaphor
 isEntailedByMetaphor
 isRelatedToMetaphorBySource
 isRelatedToMetaphorByTarget
 isSubcaseOfMetaphor
 isSourceSubcaseOfMetaphor
 isTargetSubcaseOfMetaphor
 isSupercaseOfMetaphor
 isUsedByMetaphor

Menu Help

Save

Government institution (7)

government_whole (0)

functional_status (0)

functions (0)

responsibilities (0)

position_duties (0)

sub-groups (0)

position_holders (0)

authority_structure (0)

positions (0)

position_relations (0)

authority_base (0)

Crime (5)

victim (0)

criminal (0)

criminal_activity (0)

social_problems (4)

result_of_problem (0)

society (0)

problem_cause (0)

social_problem (0)

Corruption (4)

corruption_effect (0)

corrupt_actor (0)

CORRUPTION IS A DESTRUCTIVE NATURAL PROCESS (3)

Corruption (4)

corruption_effect (0)

corruption_affectee (0)

corrupt_actor (0)

hamed_entity (0)

corrupt_activities (0)

Destructive natural process (3)

destroyed_entity (0)

destructive_process (0)

destructive_effect (0)

CORRUPTION IS DECAY (1)

Corruption (4)

corruption_affectee (0)

corrupt_activities (0)

corrupt_actor (0)

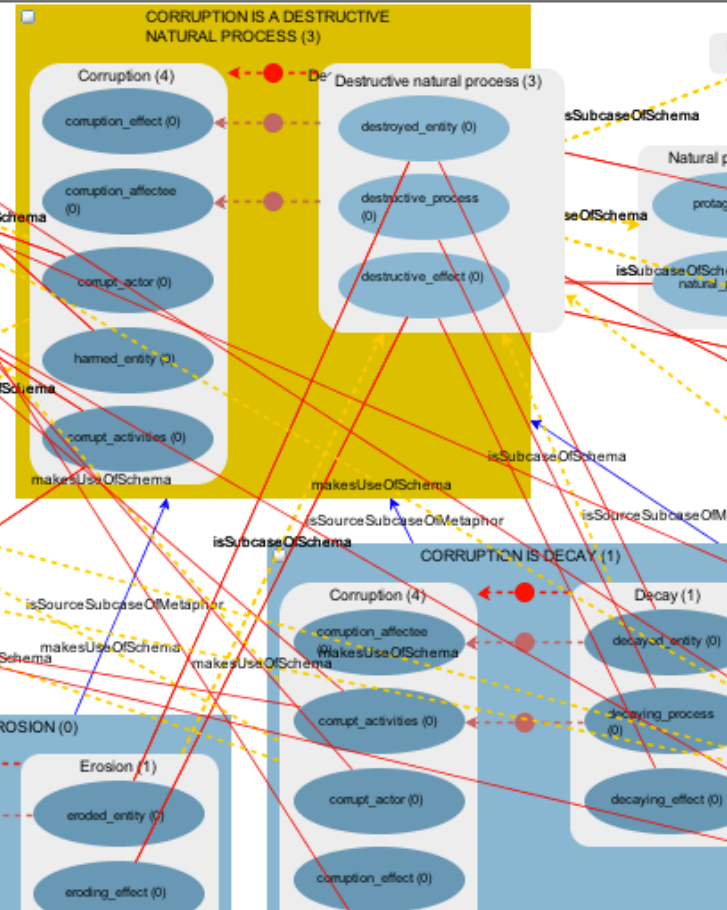
corruption_effect (0)

Decay (1)

decayed_entity (0)

decaying_process (0)


decaying_effect (0)



Validation of the cognitive aspects of metaphor

- Corpus Methods
- Basic Scalar measures
 - familiarity, accessibility, acceptability, imageability, well-formedness, conventionality, metaphoricity, informativeness, and productivity
- Behavioral Tests
 - lexical and conceptual priming, inference. measures of memorability, paraphrasing and explication, gesture, eye, body tracking
- Affective Aspects
 - Behavioral (IAT, Psychological measures)
 - Imaging
 - Metaphoric activation of emotional circuits
 - » anterior insula, and the fear and reward circuits of the amygdala and the nucleus accumbens.

Crime: Beast or Disease



Thibodeau PH, Boroditsky L (2013) Natural Language Metaphors Covertly Influence Reasoning. PLoS ONE 8(1): e52961. doi:10.1371/journal.pone.0052961

<http://www.plosone.org/article/info:doi/10.1371/journal.pone.0052961>

Conclusion

- Natural language understanding requires
 - semantic representations that support dynamic, uncertain, event based inference.
- We now understand the extraction algorithms and inference techniques that bridge multiple levels
 - NE and Extraction-based
 - Frames, Topic and Domain Models
 - Mappings and Metaphor
- Ongoing work
 - Narrative Models
 - Frame Induction
 - Metaphor Learning (Neural Computation 2013)
 - Event Synthesis (AAAI SS 2013)

Time to do this at scale!