



Lifelong Learning in Optimisation

Emma Hart Edinburgh Napier University

http://jamesobrien.tumblr.com/post/1112777561/lifel ong-learning-illustration

Optimisation Algorithms





Tuned metaheuristics Offline hyper heuristics Evolution Strategies Online hyper heuristics Metaheuristics...

Optimisation Algorithms



- Generalist: incapable of adapting to new problem characteristics
- Specialist: unable to learn from experience or exploit prior knowledge

 Contemporary ML systems usually exploit prior knowledge if faced with new but similar task



"it is now appropriate for the AI community to move beyond learning algorithms to more seriously consider systems that are capable of learning over a lifetime"

Silver, 2013

An LML should:

- Retain and/or consolidate knowledge (long-term memory)
- Selectively transfer prior knowledge when learning new tasks
- Adopt a systems approach that ensures effective and efficient interaction of elements of the system

What kind of approach might provide these features ?



An LML should:

- Retain and/or consolidate knowledge (long-term memory)
- Selectively transfer prior knowledge when learning new tasks
- Adopt a systems approach that ensures effective and efficient interaction of elements of the system

Natural Immune System

Basis of vaccination, can be very long term



An LML should:

- Retain and/or consolidate knowledge (long-term memory)
- Selectively transfer prior knowledge when learning new tasks
- Adopt a systems approach that ensures effective and efficient interaction of elements of the system

Natural Immune System Basis of vaccination, can be very long term



An LML should:

- Retain and/or consolidate knowledge (long-term memory)
- 2. Selectively transfer prior knowledge when learning new tasks
- 3. Adopt a systems approach that ensures effective and efficient interaction of elements of the system

Natural Immune System

Selectively transfer prior knowledge when learning new tasks



An LML should:

- Retain and/or consolidate knowledge (long-term memory)
- Selectively transfer prior knowledge when learning new tasks
- Adopt a systems approach that ensures effective and efficient interaction of elements of the system

Natural Immune System

Behaviour is the result of many interacting components



(Source: the Human Immune Response System www.uta.edu/chagas/images/immunSys.jpg)

The Role of the Immune Network

- Immune cells interact with other and with antigen
 - Can be stimulatory or suppressive
- Results in a network with dynamically changing topology
- Useful cells recruited into network
- Redundant ones rejected
- Topology depends on past & current environment



An LML should:

- Retain and/or consolidate knowledge (long-term memory)
- 2. Selectively transfer prior knowledge when learning new tasks
- Adopt a systems approach that ensures effective and efficient interaction of elements of the system
- 4. Generate new knowledge

Natural Immune System

Gene recombination in bone marrow continually trials news cells leading to a useful repertoire of antibodies



Immune Systems

Environment



Meta-dynamics

Computational Properties

- Exploration
 - randomly combining components from a library gives rise to many cells
- Exploitation
 - focuses search on promising cells
- Memory:
 - network provides a 'map' of the antigen space
- Adaptable
 - Doubly plastic: parametric & structural
- Diverse
 - Finite repertoire of cells has to ensure all pathogens recognised

Optimisation Systems

Environment



Meta-dynamics

Computational Properties

- Exploration
 - randomly combining components from a library gives rise to many heuristics
- Exploitation
 - focuses search on promising heuristics
- Memory:
 - network provides a 'map' of the problem space
- Adaptable
 - Doubly plastic: parametric & structural to deal with changes in problem characteristics
- Diverse
 - Finite repertoire of heuristics has to ensure all problems solved

Conceptual Overview

2d Representation of problem space





- The network sustains *heuristics* that work best in distinct regions of the instance space (diversity)
 - Need to win to be in!
- The network sustains *problems* that are representative of areas of the problem space
 - Problems that are solved by more than one heuristic are not 'interesting'
- Problems & heuristics gain concentration through mutual stimulation
 - Decay mechanisms enable gradual forgetting
 - Lack of stimulation leads to removal
- Topology of network changes over time depending on problems injected and heuristics generated

Conceptual Overview

2d Representation of problem space





- The network sustains *heuristics* that work best in distinct regions of the instance space (diversity)
 - Need to win to be in!
- The network sustains *problems* that are representative of areas of the problem space
 - Problems that are solved by more than one heuristic are not 'interesting'
- Problems & heuristics gain concentration through mutual stimulation
 - Decay mechanisms enable gradual forgetting
 - Lack of stimulation leads to removal
- Topology of network changes over time depending on problems injected and heuristics generated



Environment



Meta-dynamics

 Problem Stream
 Heuristic Generator
 Network of heuristics & problems

Environment



Meta-dynamics

1 Problem Stream

At each iteration instances *can* be injected into the system

- Single instance
- Multiple instances
- Frequent/infrequent

Environment



Meta-dynamics

- Library of components
- Components can be 'pre-defined' or evolved
- Components are combined into *heuristics*



- Library of components
- Components can be 'pre-defined' or evolved
- Components are combined into *heuristics*
- (few components -> lots of heuristics)



- Library of components
- Components can be 'pre-defined' or evolved
- Components are combined into heuristics



- Library of components
- Components can be 'pre-defined' or evolved
- Components are combined into heuristics



- Library of components
- Components can be 'pre-defined' or evolved
- Components are combined into *heuristics*



- Library of components
- Components can be 'pre-defined' or evolved
- Components are combined into *heuristics*
- Both components and heuristics can evolve



- Mutate terminal nodes
- Mutate function nodes
- Remove subtree
- Swap subtrees

- Library of components
- Components can be 'pre-defined' or evolved
- Components are combined into *heuristics*
- Both components and heuristics can evolve



Swap components Change components Remove/insert components Concatenate heuristics

- Library of components
- Components can be 'pre-defined' or evolved
- Components are combined into *heuristics*
- Both components and heuristics can evolve



- Library of components
- Components can be 'pre-defined' or evolved
- Components are combined into *heuristics*
- Both components and heuristics can evolve

Environment



Meta-dynamics



- Heuristics are stimulated by winning at least one problem
 - The higher the win, the bigger the stimulation
- Problems are stimulated if they are won by only one heuristic
 - The higher the win, the bigger the stimulation

NELLI pseudo-code

- Inject **p** problem instances with concentration *c*
- Add *h* new randomly generated heuristics with concentration *c*
 - Calculate *stimulation* of all problems P
 - Calculate stimulation of all heuristics H
 - Increment concentration of (P,H) with c
 < c_{max} and stimulation > 0
 - Decrease concentration of (H,P) if stimulation <=0
 - Remove (H,P) with c <= 0

Tune for more exploration

Depends on strength of win

Can be tuned to alter memory span & lifetime

Basic optimiser

- Bin packing:
 - 1370 instances from literature
 - All presented at start
 - Run for 500 iterations
- Number of heuristics evolved is an emergent property
- Number of problems retained gives insight into similarity of instances

	Problems solved	Extra bins
FFD	788	2142
DJD	716	2409
DJT	863	881
ADJD	686	1352
NELLI	1126	308



Generalisation Capabilities

- Train NELLI by injecting p instances and running for t iterations
- Apply evolved heuristic network to an unseen test-set:
 - Bin packing (685 train, 685 test)
 - JSSP (train, test)

Bin packing

	Problems Solved	Extra Bins
Greedy Selection	548	188
AIS Model	559	159
Island Model	557	159
NELLI	576	131

Generalisation Capabilities

- Train NELLI by injecting p instances and running for t iterations
- Apply evolved heuristic ensemble to an unseen test-set:
 - Bin packing (685 train, 685 test)
 - JSSP (train, test)

Hart & Sim (in review, J. Evolutionary Computation)

65 Taillard instances (JSSP)

	T1	Т2	Т3
EGP-JSS	0.26+/-	0.26+/-	0.26+/-
	0.04	0.03	0.010
NELLI	0.20 +/-	0.18+/-	0.18+/-
	0.09	0.03	0.04

200 new instances (JSSP)			
	T1		
Greedy Selection	68146		
GP(200P)	69795		
GP(1P)	69068		
NELLI	68125		

Memory

- a) Alternate between two different datasets every 200 iterations
- b) Present 685 randomly drawn instances each iteration



- Memory
- Learning over epoch
- Learning over lifetime

Insights into heuristic performance

- Define a **profile** (*barcode*) for each heuristic based on its relative performance on an instance
 - Results from an evolved ensemble of 8 heuristics applied to 200 unseen JSSP instances
- Could be used to directly compare diversity of heuristics using a distance metric



Insights into problem space

- Record how many heuristics 'win' each instance
- Easy problems won by many instances
- Hard problems only won by one heuristic
- Insights into relative difficulty of instances based on parameter combinations



Heuristic Perspective

• Where in the space does a particular heuristic perform well?

- Can map heuristics to problem characteristics



Summary

- NELLI (Network for lifelong Learning)
 - Generates novel heuristics that collaborate to cover instance space
 - Encapsulates memory
 - Generalises to new instances
 - Adapts to new instances
 - Better solutions than some other heuristic methods



Conclusions

- Optimisation systems should continuously learn
 - Exploit previous knowledge
 - Adapt to changing instance characteristics





Conclusions

- Optimisation systems should continuously learn
 - Exploit previous knowledge
 - Adapt to changing instance characteristics
- Optimisation systems that exploit ensembles are likely to be promising
 - No Free Lunch
 - (some) portfolio/multi-method algorithms
 - Machine learning community has plenty to say!



Conclusions

• Optimisation systems should continuously learn

- Exploit previous knowledge
- Adapt to changing instance characteristics
- Optimisation systems that exploit ensembles are likely to be promising
 - No Free Lunch
 - (some) portfolio/multi-method algorithms
 - Machine learning community has plenty to say!
- Ensemble methods like NELLI offer new insights into heuristic performance and problem difficulty
 - Create heuristic profiles that enable objective comparisons between heuristics and means of recognising diversity
 - Design better benchmarks



© 2006 Encyclopædia Britannica, Inc.



THANK YOU!





Engineering and Physical Sciences Research Council Acknowledgements Dr Kevin Sim EPSRC **EP/J021628/1**