

1 An Architecture for Adaptive Algorithmic Hybrids

2 Nicholas Cassimatis, Perrin Bignoli, Magdalena Bugjaska, Scott Dugas,
3 Unmesh Kurup, Arthi Murugesan, and Paul Bello

4 **Abstract**—We describe a cognitive architecture for creating
5 more robust intelligent systems. Our approach is to enable hybrids
6 of algorithms based on different computational formalisms to be
7 executed. The architecture is motivated by some features of human
8 cognitive architecture and the following beliefs: 1) Most existing
9 computational methods often exhibit some of the characteristics
10 desired of intelligent systems at the cost of other desired character-
11 istics and 2) a system exhibiting robust intelligence can be designed
12 by implementing hybrids of these computational methods. The
13 main obstacle to this approach is that the various relevant com-
14 putational methods are based on data structures and algorithms
15 that are difficult to integrate into one system. We describe a new
16 method of executing hybrids of algorithms using the *focus of*
17 *attention* of multiple modules. The key to this approach is the
18 following two principles: 1) Algorithms based on very different
19 computational frameworks (e.g., logical reasoning, probabilistic
20 inference, and case-based reasoning) can be implemented using
21 the same set of five *common functions* and 2) each of these common
22 functions can be executed using multiple data structures and
23 algorithms. This approach has been embodied in the Polyscheme
24 cognitive architecture. Systems based on Polyscheme in planning,
25 spatial reasoning, robotics, and information retrieval illustrate
26 that this approach to hybridizing algorithms enables qualitative
27 and measurable quantitative advances in the abilities of intelligent
28 systems.

29 **Index Terms**—Hybrid architectures, integrated systems.

30 I. INTRODUCTION

31 **W**E DESCRIBE a cognitive architecture for creating more
32 robust intelligent systems by executing hybrids of al-
33 gorithms based on different computational formalisms. There
34 are several properties that we desire of intelligent systems.
35 Each of these properties is exhibited by some algorithm, but
36 often at the cost of one of the other properties. For exam-
37 ple, many search algorithms and many probabilistic inference
38 algorithms are general insofar as a wide variety of problems
39 can be reformulated so that these algorithms can solve them.
40 They are flexible in that, when small changes are made to the
41 knowledge bases or models these methods use, they correctly
42 update their inferences. However, for larger problems, these
43 algorithms become prohibitive in time and/or space. They thus
44 trade efficiency for scalability. Reactive and behavior-based
45 systems trade generality for speed and dynamism. These sys-
46 tems can quickly adapt their behavior to new information about

the environment but are generally not capable of making many 47
kinds of complex inferences and plans that many reasoning or 48
planning algorithms can. Approaches that are based on more 49
complex and structured knowledge representation schemes, 50
such as frames [1], scripts [2], and cases [3], make the same 51
tradeoff. They can make inferences and plans that would take 52
more general algorithms too much time. However, frames, 53
scripts, and cases often do not work in situations that are even 54
slightly different from those for which they were created. 55

Because of such tradeoffs, it has been difficult to create 56
a single system that exhibits all the desired characteristics 57
of intelligent systems. One approach to this problem is to 58
create systems based on hybrids of these algorithms. This 59
paper presents a cognitive architecture, called Polyscheme, for 60
implementing and executing such hybrids. Polyscheme’s design 61
is motivated by some aspects of human cognitive architecture 62
and several computational considerations. Many methods of 63
integration involve loosely coupled collections of modules or 64
hybrids of only a small fixed set of algorithms. Although 65
Polyscheme includes modules that can encapsulate algorithms, 66
it also enables algorithms to be implemented through sequences 67
of “attention fixations” (each involving all the modules) called 68
“focus traces.” By interleaving these focus traces, Polyscheme 69
can execute hybrids of algorithms where each step of each 70
encapsulated algorithm can potentially be assisted by the oth- 71
ers. As explained hereinafter, this form of integration enables 72
systems that have previously not been feasible. 73

74 II. UNIFYING PRINCIPLES

We can motivate an architecture for integrating hybrids of 75
multiple classes of algorithms by recognizing that, even though 76
they are based on very different computational formalisms, they 77
share many common elements. 78

The following are some formal preliminaries. Strings of the 79
form $P(\dots x_i \dots, t, w)$ are called propositions. They say that 80
relation P holds over arguments x_i during temporal interval t 81
in world w . Any kind of entity can be an argument. A world 82
is a history (past, present, and future) of states. P/w refers to a 83
proposition like P except for having w as its world argument. E 84
 (“eternity”) is the temporal interval such that all other temporal 85
intervals occur during it. R is the real world. Terms can refer 86
to the same object. This is indicated with $Same(x, y, E, w)$. T , 87
 F , and U (“unknown”) are the truth values for propositions. 88

Propositions are used only to characterize certain aspects 89
of Polyscheme and as an interlingua between Polyscheme 90
modules. Polyscheme is not a “logical” system insofar, as it 91
can use nonlogical data structures and its computation is not 92
predominantly deductive or confined to manipulating formulas 93
in some logical language. 94

Manuscript received August 13, 2008; revised February 10, 2009 and June 5, 2009. This paper was recommended by Associate Editor M. Huber.

N. Cassimatis, P. Bignoli, M. Bugjaska, S. Dugas, U. Kurup, and A. Murugesan are with Rensselaer Polytechnic Institute, Troy, NY 12180 USA (e-mail: cassin@rpi.edu; bigniop@rpi.edu; bugjam@rpi.edu; dugass@rpi.edu; kurupu@rpi.edu).

P. Bello is with the Office of Naval Research, Arlington, VA 22203-1995 USA (e-mail: paul.bello@navy.mil).

Digital Object Identifier 10.1109/TSMCB.2009.2033262

95 A. Common Functions

96 It is possible to characterize algorithms from very different
97 formal frameworks as executing sequences of common func-
98 tions. In this case, “function” refers not to a mathematical
99 function but to the purpose of an algorithm as in “the function
100 of a sorting algorithm is to order the elements of a collection.”
101 These functions are described along with the notation we will
102 use to refer to them.

103 1) *Store Information*: Given the information about the truth
104 value of P , store it. $\text{Store}(P, \text{TV})$.

105 2) *Offer Opinion*: Return a truth value for a proposition.
106 $\text{OpinionOn}(P, \text{TV})$.

107 3) *Forward Inference*: Given some knowledge, infer what
108 follows. $\text{ForwardInference}(\text{BK})$, where $\text{BK} = \{\dots(P_i,$
109 $\text{TV}_i)\dots\}$ returns a list of propositions paired with their
110 respective truth values.

111 4) *Request Information/Subgoal*: In order to know about
112 something, take actions to learn about it. Given P , return a set of
113 propositions such that information about their truths will help
114 infer the truth of P . $\text{RequestInformation}(P)$.

115 5) *Identity*: For any object, find other objects to which it
116 might be identical. Where O refers to an object, BK is as
117 mentioned before, and W is an alternate world, return a set of
118 objects that might be identical to O in W . $\text{Matches}(O, \text{BK}, W)$

119 **Algorithm 1.** Gibbs Sampling

120 Represent a state variable as a proposition. Let $\{\dots P_i \dots\}$
121 be the state variables.

122 Start with an assignment of truth values to those propositions
123 $\{\dots (P_j/w_0, \text{TV}_j) \dots\}$.

124 **for** $i = 1$ to MAX-SIMULATIONS :

125 **for** each j :

126 $\text{Store}(P_j, \text{ForwardInference}(\{(P_1/w_j - 1,$
127 $\text{OpinionOn}(P_1/w_i - 1)) \dots (P_j - 1/w_i - 1,$
128 $\text{OpinionOn}(P_j - 1/w_i - 1)))$

129 $\text{Prob}(P) = \text{the proportion of worlds } w, \text{ in which } P/w$
130 is true.

131 6) *Represent Alternate Worlds*: The ability to represent and
132 make inferences about alternate states of the world is implicit
133 in all of the aforementioned common functions since they all
134 involve propositions with worlds as arguments.

135 The following are examples of how some important
136 algorithms can be implemented using common functions.

137 **Algorithm 2.** WalkSAT.

138 **for** each proposition, P , $\text{Store}(P, \text{Random}(T, F))$

139 **for** $i = 1$ to MAX-FLIPS

140 **if** $\text{OpinionOn}(C) = T$

141 Return $\{\dots (P_i, \text{OpinionOn}(P_i)) \dots\}$

142 **else**

143 with probability p ,

144 $\text{Store}(P, \text{not}(\text{OpinionOn}(R)))$

145 **otherwise**

146 $\text{Store}(\text{ForwardInference}(\text{BK}).\text{first}())$,

147 where BK is the set of ordered pairs of each
148 proposition with its truth value.

A variant of Gibbs sampling [4] can be implemented as 149
follows, where $\text{ForwardInference}(\text{BK})$ samples a value for 150
 P given the value of variables in its Markov blanket in BK . 151

Stochastic local search performed, for example, by the 152
WalkSAT algorithm [5], can be used to find solutions to 153
Boolean constraint satisfaction problems. Many problems (e.g., 154
planning, diagnosis, and circuit design) can be mapped to 155
satisfiability problems (or weighted variants of them). When 156
the following focus scheme is in place, Polyscheme performs 157
WalkSAT for a constraint C if $\text{ForwardInference}(\text{BK})$ returns 158
a single ordered pair (P, TV) , where P is the proposition such 159
that flipping its truth value will make the most clauses in C true. 160

Algorithm 3. DPLL.

161 $DPLL(P, tv)$ 162

163 $w \leftarrow$ world from the world P by assuming that the truth
value of P is tv . 164

165 **if** $\text{ForwardInference}(P/w)$ returns a contradiction

166 **return** false.

167 **if** all constraints in w are satisfied

168 **return** true.

169 $V \leftarrow \text{RequestInformation}(P/w)$.

170 **return** $DPLL(V, true) \sqcap DPLL(V, false)$.

The modern variants [6], [7] of the Davis–Putnam– 171
Logemann–Loveland (DPLL) [8] algorithm are among the 172
fastest complete satisfiability algorithms known. DPLL per- 173 AQ3
forms depth-first search through the space of assignments
of truth values to variables. As each assignment is made, 175
DPLL performs an elaboration step that makes assignments 176
that follow from the existing assignments or can be assumed 177
without contradiction. This elaboration step eliminates many 178
impossible assignments from being explored and, thus, in 179
many cases, significantly speeds search. If we assume that 180
 ForwardInference performs the DPLL elaboration step, then 181
DPLL can be reformulated as in Algorithm 3. Finally, case- 182
based reasoning (CBR) can be added to WalkSAT by modifying 183 AQ4
 ForwardInference mentioned earlier to return propositions
that were true in situations whose similarity to the current 185
situation exceeds some threshold. 186

These examples illustrate how algorithms developed within 187
very different formal frameworks can be executed with se- 188
quences of calls to the same five common functions. 189

The fact that multiple algorithms from different subfields 190
based on different formal frameworks can all be executed 191
as sequences of common functions motivates an approach to 192
integrating them. We can create hybrids of two algorithms by 193
interleaving the sequences of common functions that execute 194
each algorithm. 195

When an algorithm performs a common function on a propo- 196
sition, we say that it *attends to* or *fixes its attention* on this 197
proposition. We call this event an *attention fixation*. Each of the 198
algorithms described earlier attends to (i.e., executes a common 199
function on) one proposition at a time. 200

The sequence of attention fixations that an algorithm makes 201
when it executes is called its *focus trace*. Focus traces thus 202
provide a uniform way of characterizing the execution of 203
algorithms from different computational methods. 204

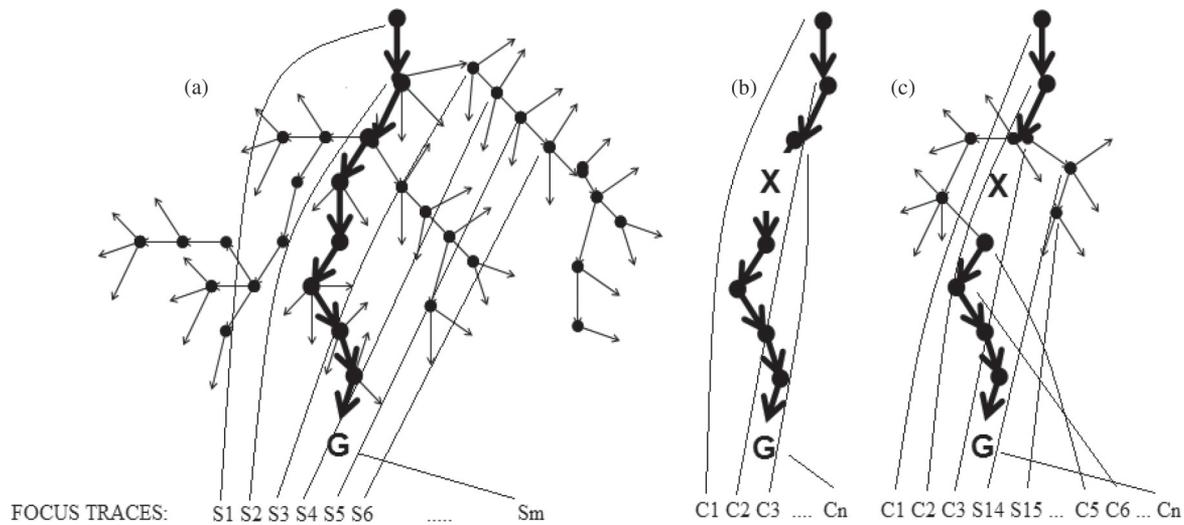


Fig. 1. Algorithms and their focus traces. The thick arrows depict a successful path to a goal found by a hypothetical run of a search algorithm. The thin arrows depict other fixations explored during the search for the best path. The (C) focus trace for the hybrid of CBR and search is a combination of the focus traces for (A) search and (B) CBR. These combine to surmount the problem (depicted by "X") by simply applying the case.

205 Interleaving the focus traces from two algorithms amounts
 206 to executing a hybrid of those algorithms. Specifically, an
 207 algorithm H is a hybrid of algorithms A1 and A2 if the focus
 208 trace of H includes fixations from A1 and A2. Fig. 1 shows
 209 the hybrid execution of CBR and search. The focus trace (1c)
 210 for the hybrid of CBR and search is a combination of the focus
 211 traces for search (1a) and CBR (1b).

212 *B. Multiple Implementation*

213 In each of the examples used to illustrate the common
 214 function principle, the common functions were implemented
 215 differently. The ability of a surprisingly diverse collection of
 216 computational methods to implement each of the common
 217 functions is what we call the *multiple implementation principle*.
 218 It is a key to enabling the integration of data structures and
 219 algorithms described herein. The following examples lend sup-
 220 port to the multiple implementation principle.

221 1) *Storing Information*: Both rule-based systems and neural
 222 networks store information. Rule matchers must keep track
 223 of which propositions they have been given as input and
 224 which they have asserted as the result of rule matches. Neural
 225 networks that implement content addressable memories store
 226 patterns using the weights of the connections between units
 227 in the network. For example, the eigenvectors of the matrix
 228 representing unit connections in Hopfield [9] neural networks,
 229 for example, correspond to patterns stored by those networks.
 230 Thus, the function of storing information is a very simple
 231 example of how very different algorithms and data structures
 232 can implement the same common function.

233 2) *Forward Inference*: Both rule-based systems and feedfor-
 234 ward neural networks, although based on very different data
 235 structures and algorithms, take inputs and produce outputs. For-
 236 ward chaining in rule-based system matches the left-hand side
 237 of a rule against a set of propositions and asserts new proposi-
 238 tions. Likewise, many forms of neural networks can be charac-
 239 terized as performing inference [10]. The values of input and

output layers of neural networks can be represented as propo- 240
 sitions [e.g., a unit representing that the temperature is 25 °C 241
 can be represented using the proposition *Temperature(25)*]. 242
 Thus, just like rule-based systems, feedforward neural networks 243
 can be characterized as taking propositions (those representing 244
 the values of the input units) and producing new propositions 245
 (those representing the values of the output units). 246

3) *Subgoaling*: Backward chaining in rule-based systems, 247
 subgoaling in logic-theorem provers, and subgoaling in means- 248
 end planners are obvious instances of the subgoaling common 249
 function. Less obviously, to learn the truth values of proposi- 250
 tions representing the output units of neural networks, one can 251
 perform a kind of subgoaling by finding (the truth values of 252
 propositions representing) the values of the input units. Even 253
 directing visual attention is a form of subgoaling. In this case, 254
 to learn about the truth value of a proposition representing the 255
 attribute of an object, a goal is made of directing a camera or a 256
 pair of eyes toward that object. 257

4) *Identity Matching*: Algorithms used in object recognition 258
 (e.g., Bayesian networks, neural networks, and support-vector 259
 machines) all are performing a kind of similarity matching by 260
 taking a visually perceived object and finding the most similar 261
 stored object. Case-indexing and retrieval schemes in CBR have 262
 the same function. 263

5) *Representing Alternate Worlds*: Mechanisms which can 264
 make inferences about the real world can make inferences about 265
 hypothetical worlds. 266

These examples illustrate that common functions can be 267
 implemented using computational methods from very different 268
 subfields of artificial intelligence. 269

Implementing a single common function with multiple dif- 270
 ferent algorithms and data structures enables another kind of 271
 hybridization. Consider an algorithm A executed in a particular 272
 way using common functions C1...Cn. If each of the C is 273
 executed using multiple computational methods M1...Mn, then 274
 every step of the execution of algorithm A will also involve the 275
 methods M. 276

277 As another example, one could imagine Gibbs sampling
 278 executed using a sequence of common functions, each of which
 279 is executed by a neural network. Thus, although the neural
 280 network would be performing all the computations (in this case,
 281 computing the likelihood of a proposition given its Markov
 282 blanket), the focus of Polyscheme can be selected in such a way
 283 that it implements Gibbs sampling. If one of the M_i executing
 284 a common function is an algorithm processing new sensor
 285 information, then every step of inference executing using this
 286 common function would incorporate up-to-date information
 287 from the world.

288 C. Cognitive Self-Regulation

289 The common function principle shows that very different
 290 kinds of algorithms can be executed as sequences of the same
 291 set of common functions. How does an intelligent system
 292 implementing more than one algorithm decide which common
 293 functions to choose? A recurrent theme among the algorithms
 294 used to illustrate the common function principle is that they
 295 choose which common function to execute in response to
 296 *metacognitive problems*. This is evident in several broad classes
 297 of algorithms.

298 Search algorithms choose possible worlds by assuming the
 299 truth of propositions that are *unknown*. If the truth value of a
 300 proposition is known, it is considered fixed, and worlds with
 301 that proposition having the opposite value are not explored. In
 302 the case of constraint-based search algorithms, the *metacogni-*
 303 *tive* problem is *ignorance*, which occurs when the truth value
 304 of a proposition is not known. In the case of search-based
 305 planning, the ignorance is often based on *conflict* since there
 306 is more than one possible next action to take or explore.

307 Many stochastic simulation algorithms for probabilistic rea-
 308 soning (e.g., Gibbs sampling) also choose worlds to explore
 309 based on unknown values of a state variable. However, the
 310 level of ignorance in this case is somewhat reduced since a
 311 probability distribution for the state variable is known. This
 312 leads to a somewhat different strategy of exploring worlds:
 313 Worlds with more likely propositions being true are likely to
 314 be explored more often.

315 CBR algorithms are also often driven by ignorance and differ
 316 from search and stochastic simulation algorithms by how they
 317 react to ignorance. Instead of exploring possible worlds where
 318 different actions are taken, they retrieve similar solutions to
 319 similar problems and try to adapt them.

320 Thus, CBR, search, and stochastic simulation are different
 321 ways of addressing different kinds of *metacognitive problems*.
 322 We call the insight that many algorithms can be character-
 323 ized by which common functions they choose in response to
 324 metacognitive problems the *cognitive self-regulation principle*.

325 This analysis resembles Soar's [11] implementation of uni-
 326 versal weak methods by reacting to impasses (analogous to
 327 metacognitive problems) by reasoning in problem spaces (anal-
 328 ogous to alternate worlds). However, together with the common
 329 function and multiple implementation principles, the cognitive
 330 self-regulation principle motivates a significantly different ar-
 331 chitecture.

III. POLYScheme

332

The principles¹ from the last section motivate the design 333
 of a cognitive architecture called Polyscheme. In Polyscheme, 334
 algorithms can be implemented as strategies for focusing the 335
 attention of multiple modules. By combining strategies, we 336
 lay down hybrid focus traces (as described in the last section) 337
 and thus execute hybrids of algorithms implemented as focus 338
 control strategies. By using modules based on different data 339
 structures and algorithms, we enable hybrids of algorithms 340
 implemented as focus control strategies and the algorithms in 341
 the modules. By including modules that can sense the world 342
 and by tightly bounding the time each focus of attention takes, 343
 inference can constantly use and react to the latest information 344
 from the world. This section and the next elaborate these points. 345

A Polyscheme system is comprised of a set of specialist S, 346
 an attention buffer A, and a focus manager FM. FM implements a 347
 getNextFocus() procedure that returns the next proposition to 348
 attend to. All the specialists must implement the following pro- 349
 cedures: Store(), OpinionOn(), Matches(), and ModifyFM(). 350
 The latter influences which propositions FM selects. It will be 351
 described in more details later in this section. 352

At every time step, FM chooses a proposition for all the 353
 specialists to "focus on." Specifically, specialists give their 354
 opinions on the proposition's truth value, get all the other 355
 specialist's opinions on it, make their own inferences based 356
 on this new information, and request other propositions for the 357
 FM to focus on. The reasons why all specialists focus on the 358
 same proposition at the same time are as follows: 1) so that no 359
 specialist makes inferences based on a truth value that another 360
 specialist knows to be incorrect and 2) because, even though a 361
 proposition is focused on because specialist S1 requested to fo- 362
 cus on it, it might become relevant for some other specialist S2. 363

More formally, the Polyscheme control loop is 364

Do forever :

"Select the next focus."

$P = FM.getNextFocus()$.

"Get each specialists opinion on the focal prop."

For all specialists S :

$TV_i = S.opinionOn(P)$

"Inform specialists of each other's opinions."

For all specialists S :

For each of the TV_i

$S.Store(P, TV_i)$

"Allow specialists to influence the Focus Manager."

For the focus queue Q and all specialists S :

$S.modifyFM(P, Q)$

Once a focus of attention is chosen, all the specialists 365
 must focus on and therefore execute the same functions on it. 366

¹There are also several aspects of human psychology that motivate the main aspects of Polyscheme. These are described in [12] and [13].

367 Therefore, the choice of attention fixation of the FM controls the
 368 flow of computation. We have experimented with several kinds
 369 of FMs but will use a very simple FM based on a queue (FQ)
 370 to illustrate how attention selection can implement hybrids of
 371 many kinds of algorithms. Although achieving all our long-term
 372 objectives will almost certainly require a more sophisticated fo-
 373 cus manager, a queue-based focus manager has been sufficient
 374 to achieve significant results and demonstrate the promise of the
 375 approach. Thus, in what follows, `ModifyFM()` takes a queue as
 376 input and modifies it. These modifications can involve adding
 377 elements, deleting them, and changing their order.

378 The principal means of interaction among specialists is the
 379 sharing of opinions in Polyscheme’s control loop. All special-
 380 ists are first asked their opinion on the focal proposition, and
 381 then, each learns about these opinions and processes them.
 382 Therefore, specialists must wait for all of other specialists to
 383 finish offering their opinion before they proceed to use these
 384 opinions. This keeps specialists from making an inference on
 385 the opinion of one specialist that is contradicted by another
 386 specialist during the same time step.² If a proposition’s truth
 387 value is inferred by a specialist during one iteration of the loop,
 388 other specialists will only learn about it if that proposition is
 389 focused on at a subsequent time step.

390 To illustrate, consider a Polyscheme system that has a rule
 391 specialist with $A \rightarrow B$, a perception specialist that can see that
 392 A and C are true, and a neural network specialist that classifies
 393 B and C as a situation represented by the proposition D . This
 394 licenses the following inferences: B is true (because A is true),
 395 and thus, D is true (because B and C are classified as a situation
 396 where D is true). The following is a sketch of how information
 397 would flow through the focus of attention in Polyscheme to
 398 generate these inferences.

- 399 1) The perception specialists puts A and C on the focus
 400 queue.
 401 a) Summary: The perception specialist requests that spe-
 402 cialists focus on what it has seen to be true.
 403 2) At a later iteration of the control loop, A is chosen for
 404 focus.
 405 a) Summary: A is focused on, and perception specialist
 406 says it is true; rule specialist infers that B is true and
 407 requests for focus on B .
 408 b) Taking opinions
 409 i. The perception specialist asserts its opinion that A
 410 is true.
 411 c) Storing opinions
 412 i. The rule specialist infers that B is true.
 413 d) Requesting focus
 414 i. The rule specialist puts B on the focus queue.

- 3) At a later iteration, B is chosen for focus. 415
 a) Summary: “ B is focused on, rule specialist asserts it 416
 is true (because of the rule $A \rightarrow B$), and the neural 417
 network specialist commits this to memory.” 418
 b) Taking opinions 419
 i. The rule specialist asserts its opinion that B is true. 420
 c) Storing opinion 421
 i. The neural network specialist stores in its memory 422
 that B is true. 423
 d) Requesting focus 424
 i. None of the specialists infers anything new; thus, the 425
 focus queue is not changed. 426
 4) At a later iteration of the control loop, C is chosen for 427
 focus. 428
 a) Summary: “ C is focused on, and perception specialist 429
 asserts it is true; neural network classifies this as 430
 situation D and requests that D be focused on.” 431
 b) Taking opinions 432
 i. The perception specialist asserts its opinion that C 433
 is true. 434
 c) Storing opinions 435
 i. The neural network specialist stores in its memory 436
 that C is true and classifies this as a situation D . 437
 d) Requesting focus 438
 i. The neural network specialist requests focus on D . 439
 5) At a later iteration, D is chosen for focus. 440
 a) Summary: “ D is focused on, the neural network spe- 441
 cialist asserts it is true, and all the other specialists 442
 learn this.” 443
 b) Taking opinions 444
 i. The neural network specialist asserts its opinion that 445
 D is true. 446
 c) Storing opinions 447
 i. All the specialists learn that D is true. 448

This example illustrates that, if a specialist makes an infer- 449
 ence about proposition P at one time step, it is only learned by 450
 other specialists after P is focused on at a future time step. This 451
 is a relatively cumbersome process for making what intuitively 452
 appears to be a simple two-step inference. However, as we 453
 will demonstrate hereinafter, implementing inference through 454
 the focus of attention of multiple specialists can thus generate 455
 significant benefits. 456

IV. ALGORITHMS AS FOCUS CONTROL STRATEGIES 457

In this section, we show how to implement algorithms using 458
 focus management in Polyscheme and illustrate the integration 459
 that this enables. We will concentrate mostly on algorithms 460
 from Section II. In that section, algorithms were described 461
 using common functions. In this section, specialists imple- 462
 ment those common functions. We illustrate how focus control 463
 chooses which propositions those common functions operate 464
 on. In each case, a key component of the implementation is how 465
 the `ModifyFM` procedure is implemented. In the description of 466
 this procedure for each algorithm, we use italics to empha- 467
 size the role of a metacognitive problem (such as uncertainty, 468
 conflict, or ignorance) in choosing the focus of attention. This 469
 illustrates how cognitive self-regulation (through the choice of 470

²This implies that, if a specialist finishes computing its opinion before
 other specialists, it must waste time waiting for them to finish. In practice,
 two factors mitigate this waste. First, when a system is run on one or a few
 processors, idle CPU time is allocated to specialists not done computing,
 and there is little waste. Second, specialists are often designed to return
 their opinions very quickly. This minimizes the time they spend waiting for
 other specialist to finish. Such waste as there is, however, is compensated for
 significantly by the fact that specialists will never make inferences based on
 information another specialist knows to be false. The results reported later in
 this paper demonstrate that, despite this waste, we can still achieve significant
 performance improvements.

471 attention fixation) is a common aspect of the flow of control in
472 the algorithms we describe.

473 First, let us consider how to implement a local-search
474 algorithm such as WalkSAT. We can implement pure WalkSAT
475 using a “constraint specialist.” At any given time, this specialist
476 encodes a constraint C in a conjunctive normal form. The
477 specialist’s functions operate as follows:

Store(P, TV) : P and TV are added as a clause to C .

478 ModifyFM(P, Q): If C is satisfied in the world w of P , then
479 add Satisfied(C, E, w) to Q . Otherwise, put the proposition
480 P involved in most clauses with *conflicting truth values* at the
481 front of the focus Q

OpinionOn(P) : If P is of the form

Satisfied(C, w), return T if C is satisfied in w ;

otherwise, return U.

482 DPLL can be executed by a different kind of constraint
483 specialist. Storing a proposition performs elaboration

Store(P, TV) : P and TV are added as a clause to C ;

perform the DPLL elaboration step on C .

484 ModifyFM(P, Q): If C is satisfied in the world w of P , then
485 add Satisfied(C, E, w) to Q . Otherwise, randomly choose a
486 proposition P whose truth value remains uncertain and put it,
487 and its negation, at the front of Q . (It is common to improve
488 DPLL with variable selection heuristics. There is no obstacle to
489 incorporating these heuristics into ModifyFM.)

OpinionOn(P) : If P is of the form

Satisfied(C, w), return T if C is satisfied in w ;

otherwise, return U.

490 A system with repeated experience in similar environments
491 will have made many inferences and solved many problems.
492 One can speed it up using a form of CBR. A “case specialist”
493 remembers the previous states of the worlds and the relations
494 among objects in those situations. In a new situation, it is
495 capable of finding similarities with the past and suggesting
496 cases relevant to the present.

497 More specifically

Store(P, TV). (P, TV) is added to memory.

498 ModifyFM(P, Q). If the *truth value of P is uncertain* and it is
499 involved in a structure (i.e., a set of related propositions) that
500 is highly similar (above some threshold) to a proposition P_1 in
501 a previously encountered structure, then transform the old case
502 into propositions that use the objects and times in the current
503 situation, and put those propositions on the queue.

504 The three questions raised by this approach to CBR involve
505 how the relevant features to a case are focused on, how cases
506 are stored in memory, and how similarity between cases is
507 computed. The last two questions are beyond the scope of this
508 paper because our goal here is to provide a framework for

integrating computational methods and not determining which 509
specific methods are worth integrating. As for the first question, 510
some of the features of a case must be focused on because they 511
were perceived or inferred by other specialists. These will cue 512
a case. The ModifyFM procedure described earlier would then 513
have the result of adding focus for the rest of the features. 514

As an example of integration using focus traces, consider 515
that CBR and WalkSAT integrate easily. WalkSAT proceeds as 516
it normally would, flipping the truth value of one proposition 517
at a time and focusing on it. When the case specialist finds a 518
situation in the past that is sufficiently similar to this one, it 519
essentially flips several truth values at once, hopefully getting 520
WalkSAT much closer to a solution. The brittleness of CBR is 521
ameliorated by the fact that whatever inconsistencies there are 522
between the past case and the current situation can be resolved 523
by continued search by WalkSAT. 524

Finally, for environments that are changing and/or are per- 525
ceptible only through noisy sensors, every step of CBR and 526
WalkSAT should be influenced by the information from these 527
sensors. 528

Adding a “perception specialist” achieves this 529

ModifyFM(P, Q). If sensors perceive that a proposition

X has *changed its truth value*, put X

at the beginning of Q .

OpinionOn(P). Return the truth value that P had

the last time it perceived it or

U if you have never perceived it.

Thus, during any attention fixation caused by the constraint 530
or case specialist, if new information about a proposition P is 531
perceived, the perception specialist will ask the focus manager to 532
focus on it. When Polyscheme does focus on P , the perception 533
specialist will take the appropriate stance on it. The case and 534
constraint specialists will then learn about it through their 535
Store() procedures and incorporate the information into their 536
case matching and search, respectively. 537

These examples illustrate that algorithms from very different 538
computational frameworks can be integrated using the same 539
“computational building blocks,” i.e., the focus of attention of 540
a set of modules. 541

V. BENEFITS OF FOCUS TRACE INTEGRATION 542

We will describe how Polyscheme enables the best char- 543
acteristics of multiple diverse computational methods to be 544
combined by describing a working mobile robot [14] controlled 545
by Polyscheme. The robot’s goal (shown in Fig. 2) was to keep 546
track of and follow another robot as it moved about and occa- 547
sionally became occluded by other objects. This task requires 548
many of the characteristics of intelligent systems we described 549
in the introduction. The reasoning and planning problems were 550
difficult. Because the actions of the robots had side effects and 551
because there was incomplete information (because of occlu- 552
sion), traditional planners could not be used. Formulating the 553
physical constraints on objects (e.g., they do not pass through 554

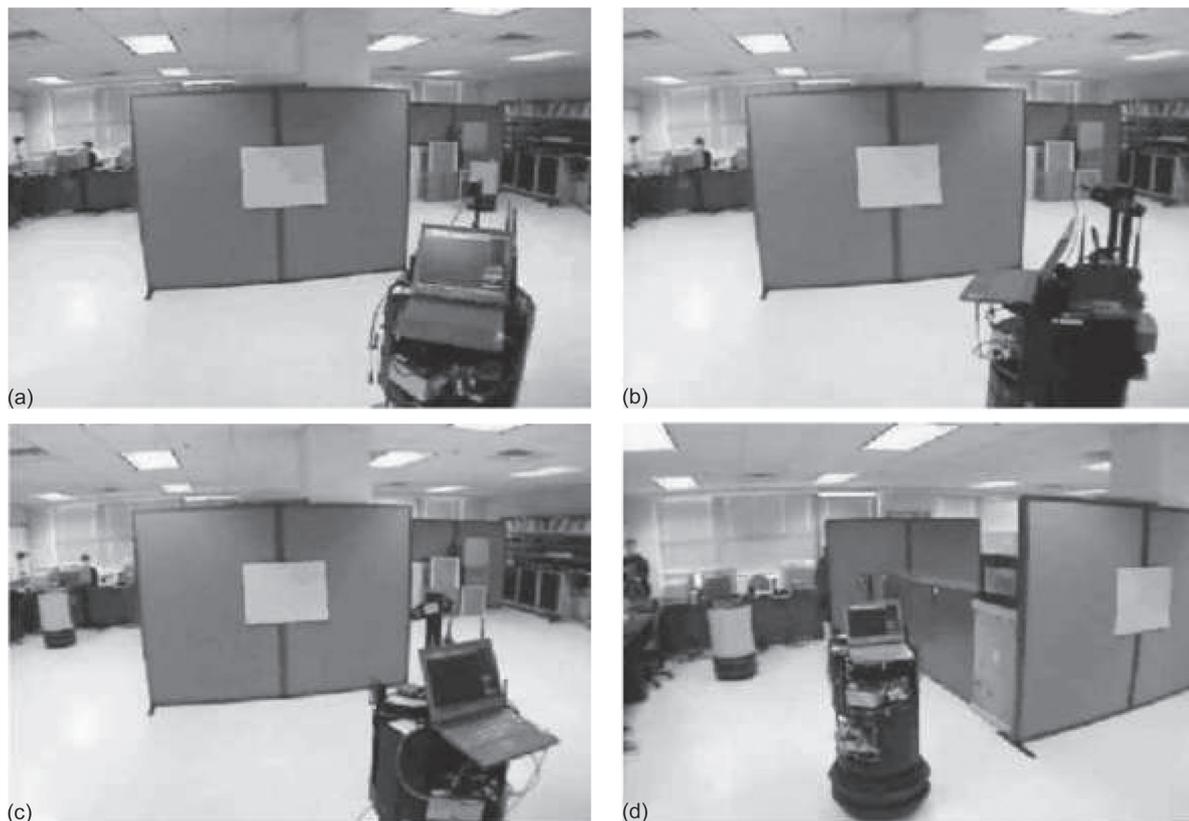


Fig. 2. Robot tracking problem. The robot in the foreground of (a) is tasked to track the robot in the distant right of (a). The tracked robot (b) disappears behind an occluder, and a robot with the exact appearance (c) emerges from the left of the occluder. Polyscheme infers that the two are identical and moves the tracking robot toward the visible robot. When (d) it sees an obstacle to robot motion behind the occluder, it infers that the two robots are different.

555 each other, gravity, etc.) in this domain in a SAT solver was
 556 prohibitive. Memory demands grew exponentially with spatial
 557 resolution, exhausting the memory of typical desktop comput-
 558 ers with grid sizes around $5 \times 5 \times 5$. Furthermore, although
 559 operating in an environment while taking into account physical
 560 laws and potentially occluded objects is well beyond the reach
 561 of pure reactive systems, the positive characteristics of these
 562 systems were required. The environment was constantly chang-
 563 ing; thus, the robot had incomplete information, and it therefore
 564 needed to be reactive and flexible enough to incorporate new
 565 information from its sensors into its planning and inference.

566 We briefly describe the Polyscheme system that controlled
 567 the robot. The focus manager was a modification of the queue
 568 scheme described earlier. The main modification was that
 569 propositions in the queue were associated with “satisfaction
 570 conditions” that would cause the proposition to be removed
 571 from the queue when they were met. For example, if proposition
 572 P1 was put on the queue to help infer if P2 was true, then if the
 573 truth of P2 is determined, there is no longer a need to focus
 574 on P1, and it is removed. The specialists included a perception
 575 specialist that detected the location and type of objects in
 576 the environment. A physical constraint specialist kept track of
 577 physical constraints, and a path specialist included a library of
 578 “path scripts” that described paths that robots typically take.

579 We can now use this robot to help describe how Polyscheme
 580 enables systems that exhibit the characteristics of algorithms
 581 based on diverse computational formalisms.

A) *Generality and Flexibility:* Methods such as local search, 582
 backtracking search, and stochastic simulation make infer- 583
 ences, find plans, or solve constraints in a very wide variety 584
 of domains. This makes them general and flexible in that, when 585
 a situation changes and is formulated for these algorithms, they 586
 will deal with them accordingly. We have already described 587
 how to implement such algorithms within Polyscheme. In our 588
 robot, these kinds of algorithms were used to maintain the 589
 physical constraints described before. The benefit of doing so 590
 is that the system can also exhibit the following characteristics, 591
 which are not often exhibited in pure versions of these general 592
 algorithms when operating on many classes of problems. 593

B) *Speed:* General algorithms tend to be slow on larger 594
 problems because the state spaces they explore grow very 595
 quickly as a problem grows. “Structured” reasoning and plan- 596
 ning algorithms based on frames or scripts do not have this 597
 problem since they do not generally search state spaces but in- 598
 stead make inferences or solve problems using large structured 599
 representations (i.e., frames, scripts, or cases). When these 600
 algorithms are implemented in Polyscheme together with more 601
 general and flexible algorithms, the best characteristics of each 602
 can be exhibited in the same system. In our robot, Polyscheme 603
 could find a continuous path between two sightings of a tracked 604
 object more quickly than pure search because of its library of 605
 path scripts. When a script was retrieved that did not completely 606
 match the existing situation, the constraint system would detect 607
 this contradiction, and this would initiate a search for a model 608

609 for changes to the script that would be more consistent with
610 the situation. While full search was always available, the script
611 retrieval effectively meant that the search began in a state much
612 closer to the correct model of the world. Thus, the speed of
613 structured approaches was combined with the generality and
614 flexibility of search-based methods.

615 *C) Reactivity:* Environments whose states change and
616 which are sensed through imperfect sensors require systems to
617 be able to quickly update their plans and inferences upon new
618 information. Reactivity can be achieved using Polyscheme by
619 mandating each focus of attention to be quick and by including
620 sensor specialists. This guarantees that little inference will hap-
621 pen before new information is sensed. When our robot moved
622 to a different location and saw objects that were formerly
623 occluded, the perception specialist requested that all the other
624 specialists focus on this information immediately. When they
625 did, the constraint and path specialists were able to update their
626 inferences and plans accordingly.

627 The ability of the robot to engage in complex reasoning and
628 planning while moving about in a changing world demonstrates
629 that creating hybrids of algorithms in Polyscheme can enable a
630 combination of characteristics not possible given the existing
631 individual methods alone.

632 VI. EVALUATIONS AND IMPLEMENTED SYSTEMS

633 We used Polyscheme to create several systems that illustrate
634 the benefits of this approach to hybrid algorithm execution.
635 These benefits can be measured quantitatively, and they can be
636 illustrated by systems that provide qualitatively new function-
637 ality.

638 A. Quantitative Evaluations

639 To confirm that Polyscheme’s integration of multiple data
640 structures and algorithms to constrain inference could lead to
641 computational speedups, we performed quantitative evaluations
642 on path-planning and spatial-reasoning problems.

643 *1) Path Planning:* Finding a path through a graph is an
644 important problem in several fields, particularly robot motion
645 planning. It is common to discretize a continuous space into
646 a graph and use an algorithm such as A^* [15] to find the
647 optimal path through the graph. This approach requires several
648 assumptions to be made that cannot be obtained in many real-
649 world environments. The most important of these assumptions
650 involves change. For example, these algorithms typically do not
651 enable planning to account for a change in a robot’s abilities.
652 If a robot that can fit through a door picks up and carries a
653 large object, it may no longer be able to move through the
654 door. Thus, the link between the regions that the door connects
655 must be severed in the graph. However, most path-planning
656 algorithms assume a fixed graph. Another example of change
657 that traditional motion planning algorithms do not account for
658 involves changes in the environment that do not involve the
659 robot. For example, if a door closes automatically during certain
660 temporal intervals, then the graph representing the connectivity
661 of the regions must change. (Even path-planning algorithms,
662 such as Ariadne [16], that use multiple kinds of algorithms to

improve search do not deal with change through time of this
663 sort.) For reasons such as these, dedicated motion planning
664 algorithms cannot work in many real-world situations. 665

666 One solution to addressing the problem of action effects
667 and change is to formulate path planning as a weighted SAT
668 problem. Side effects and changes of state would be easy to
669 formulate as SAT constraints. However, to represent that the
670 states of objects can change over time, SAT encodings of
671 such problems must include a copy of each state variable for
672 every time step. For example, it is not sufficient to have a
673 $door17Open$ variable. One must have $door17OpenAtTime1$,
674 $door17OpenAtTime2$, etc., variables. Thus, the size of the
675 required SAT formulation would be very large for situations
676 that extend over many time steps. However, the execution time
677 of SAT algorithms, as will be illustrated hereinafter, grows
678 exponentially with the number of variables. Thus, for problems
679 with many time steps, SAT encodings are not an efficient means
680 of planning paths.

681 One solution to this problem is to introduce temporal inter-
682 vals. Rather than write that $door17$ is open at times 3, 4, 5,
683 6, 7, 8, 9, 10, and 11, one can write that $door17$ is open over
684 the interval (3, 11). However, SAT encodings would require
685 that every possible interval be represented. Thus, for example,
686 a domain with 100 time points entails tens of thousands of
687 temporal intervals with those times as end points. Furthermore,
688 this formulation would not reduce the search space. To move
689 from A to B , the search algorithm would have to consider the
690 world where the motion occurred at time 1, the world where
691 it occurred at time 2, and so on, even though, in most cases,
692 the specific time is irrelevant. The root cause of the problem is
693 that the SAT formulation requires all constraints to be grounded
694 propositionally and that it cannot reason over indefinite objects.
695 If they could, then they could plan under the assumption, e.g.,
696 that the robot moved from A to B at an “indefinite” time t and
697 the only reason about the specific bounds on t if they become
698 relevant. 699

700 To enable reasoning over new objects, we created a system
701 for “generative” SAT (GenSAT) solving called GenDPLL [17].
702 It implements DPLL in the manner outlined in Section III.
703 However, rather than encoding constraints as propositional SAT
704 problems, it encoded them using first-order constraints. These
705 constraints were stored in a “formula specialist” that used a rule
706 matching algorithm to perform the DPLL elaboration step. 707

708 This enables intervals to be represented as “indefinite” times
709 that are constrained between two initially unknown points. For
710 example, the following constraint encodes that, if an object
711 moves from being in A at time a to being in nonadjacent B
712 at b , it must have traveled to an intermediate adjacent location
713 at some time point 714

$$715 \begin{aligned} & Loc(?x, ?p_1, ?t_1) \wedge Loc(?x, ?p_2, ?t_2) \wedge \neg Same(?p_1, ?p_2) \\ & \rightarrow Adjacent(?p_1, ?p_x) \wedge Meets(?t_1, ?t_x), ?p_x, ?t_x. \end{aligned}$$

716 Thus, indefinite times eliminate the need to consider all the
717 possible end points of t , significantly reducing the search space. 718

719 We tested the increasing efficiency so enabled in a robot
720 path-planning domain. We presented Polyscheme and planners
721 based on modern weighted SAT solvers a 2-D grid and two end 722

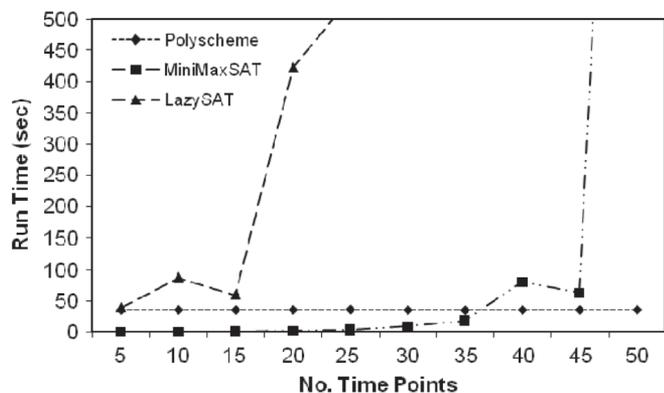


Fig. 3. Performance of Polyscheme, MiniMaxSat, and LazySat on a path-planning problem.

717 points. The grid contained obstacles whose locations changed
718 over time and, thus, for the reason mentioned before, required
719 the capability to reason about change. The goal of these systems
720 was to find the shortest possible path between points. The fact
721 that objects could change properties and locations made many
722 common path planners unusable for this domain and, for the
723 reasons mentioned earlier in the section, made this a very hard
724 search problem for conventional SAT algorithms. Because we
725 desired the ability to include constraints that were not purely
726 about paths, we could not simply consider time as a third
727 dimension and use a 3-D path planner.

728 We compared Polyscheme’s performance against LazySAT
729 [18] and MiniMaxSAT [19]. LazySAT was used because it
730 is the only weighted MaxSAT solver that lazily instantiates
731 constraints. LazySAT uses a WalkSAT-like [5] local-search
732 algorithm. Since such algorithms, in many cases, perform worse
733 than systematic solvers, we also evaluated MiniMaxSAT’s per-
734 formance. MiniMaxSAT is a weighted SAT solver based on
735 MiniSAT, which is, in turn, an extension of MiniSAT. MiniSAT,
736 the winner of the 2006 SAT Competition, is widely regarded as
737 one of the best available SAT solvers.

738 As Fig. 3 shows, the performance of GenDPLL was approx-
739 imately constant in the number of time points and performed
740 faster than LazySAT on problems with more than approx-
741 imately five time points and MiniMaxSAT with more than
742 35 time points. The graph reflects CPU times and thus demon-
743 strates that, even though Polyscheme involves significantly
744 more computation for any step in the search, its ability to use
745 multiple data structures and algorithms to prune search is often
746 able to significantly compensate for this overhead.

747 Averages of ten runs for each number of time steps were used
748 in forming the graph in order to average away changes in per-
749 formance due to background processes on our computer or the
750 specific problem instances. Since LazySAT is not guaranteed
751 to halt with an optimal solution, the times displayed indicate
752 the speed with which it found a path that did not break hard
753 constraints (e.g., that robots cannot pass through solid objects).
754 Although LazySAT lazily instantiates constraints, it does
755 not reason over unknown objects and thus achieves neither
756 reductions in the number of explored models nor the consequent
757 speedups that are afforded by the GenSAT interval formulation.
758 In the indefinite interval formulation, merely adding time

steps does not change the number of models that need to be
759 considered. 760

761 These results demonstrate that, by using Polyscheme to
762 create a hybrid of a SAT-solving algorithm (GenDPLL) and
763 a rule matcher, problems which were not solvable by these or
764 other methods alone can be solved efficiently. 765

766 2) *Spatial Reasoning*: Domains that involve spatial rela-
767 tions can pose a problem for inference algorithms because of
768 the number of points they involve. A 100×100 grid, for
769 example, contains 10000 points. If one is uncertain about the
770 location of A and only knows that point B is within ten cells
771 away from B , then there are on the order of one million possible
772 configurations of A and B in that grid that are consistent with
773 this knowledge. 774

775 There are many “diagrammatic” reasoning systems that en-
776 able reasoning with such spatial constraints. They tend to be
777 more efficient than, for example, SAT solvers on such problems
778 in part because diagrams more compactly represent spatial re-
779 lations. Existing diagrammatic reasoning systems are however
780 quite limited in their ability to also reason over nonspatial con-
781 straints, and no such systems offer the generality, soundness,
782 and completeness of SAT solvers. 783

784 To address this problem, we created a hybrid system in-
785 tended to provide the benefits of SAT methods while enabling
786 considerably more efficient inference on problems with spatial
787 relations. The system was able to take input in the same form as
788 other SAT solvers, except that some of the constraints could
789 involve (possibly metric) spatial relations. Examples include
790 $Near(a, b, 10)$ (“ a is within ten units of b ”), $Left(b, c)$, and
791 $Above(a, c)$. Constraints could mix spatial and nonspatial re-
792 lations so that it would be possible to represent a constraints to
793 the effect of “Dogs on a leash are near the person holding the
794 leash.” Given such constraints as input, the system finds models
795 that satisfy them, if they exist. 796

797 This system, called DPLL-S, was based on a weighted SAT
798 solver similar to the one described in the last section. In
799 addition, the system included a diagram component that kept
800 track of spatial relations and used “possibility spaces” [20]
801 to compactly represent relations. Thus, rather than needing
802 to represent every specific possible location of an object, it
803 could represent the region (i.e., possibility space) where the
804 object could potentially be located. This representation would
805 enable the diagram component to rule out many possible object
806 locations very quickly. For example, if A is more than 70 units
807 from B and C is within ten units of B , then, using standard
808 grid algorithms, one can very quickly infer using possibility
809 spaces that A is not near C . The only way to rule out that
810 possibility using a standard SAT solver would be to search over
811 potentially millions of possible configurations of A , B , and C .
812 The diagram specialist can therefore dramatically reduce the
813 number of possibilities explored during search. 814

815 To measure the impact of this hybridized approach, we
816 tested the system on problem instances of varying size. Since
817 a criticism of many diagrammatic reasoning systems has been
818 that they are tested only on problems that suit them, we
819 randomly generated problems. We then solved these prob-
820 lems using MiniMaxSAT and our hybrid system. On problems
821 with smaller grid sizes, MiniMaxSAT outperformed the hybrid
822 816

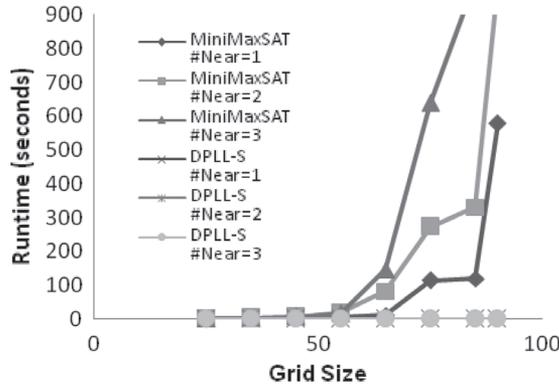


Fig. 4. Performance of DPLL-S versus MiniMaxSAT.

817 system. However, with increasing grid sizes, the benefits of
 818 reasoning with possibility spaces increased, and the hybrid sys-
 819 tem significantly surpassed MiniMaxSAT’s performance. Fig. 4
 820 shows these results. “*Near = n*” in the legend indicates that
 821 the problems run had the no *Near* predicate had an argument
 822 greater than *n*. As *n* grows, there are generally more possible
 823 configurations that satisfy a constraint, and thus, performance
 824 deteriorated.

825 3) *Computational Complexity of Search*: The path-planning
 826 and spatial-reasoning problems we have been discussing have
 827 a high degree of computational complexity. We make no claim
 828 that our approach somehow changes the complexity class of
 829 these problems nor do we claim to have surmounted the “no-
 830 free lunch” [21] theorems, which hold that “any two algo-
 831 rithms are equivalent when their performances are averaged
 832 across all possible problems” [22]. It is however very common
 833 for algorithmic improvements that do not change complexity
 834 characteristics to nevertheless make an approach tractable on
 835 problems that had been impossible before. Innovations such
 836 as stochastic local search [23] and clause learning [6] have,
 837 for example, made SAT solvers usable on a wide range of
 838 new problems without altering the fact that SAT solving is a
 839 nondeterministic polynomial time complete problem. Although
 840 the technical details are quite different, Polyscheme follows in
 841 this tradition.

842 Furthermore, the performance gains mentioned here were
 843 specific to domains with spatial and temporal relations. This
 844 is to be expected since the gains were achieved by hybridizing
 845 search with specialized spatial and temporal reasoning meth-
 846 ods. While the specificity of these improvements is a limitation
 847 of these systems, there is a wide body of research in cognitive
 848 science [24], indicating that reasoning in many domains can be
 849 mapped onto reasoning about a set of relations such as time and
 850 space. Confirming the potential breadth of application of such
 851 specializations is a topic for future research.

852 Finally, Polyscheme involves a significant amount of over-
 853 head that requires considerably more computations per state
 854 explored during search. This is confirmed by our quantitative
 855 evaluations. However, in the case of path planning, we showed
 856 that the new problem formulation that Polyscheme enables
 857 significantly reduces the search space and, hence, the CPU
 858 time needed, as the problems involved larger time frames. The
 859 spatial-reasoning search showed how a diagram module could

significantly reduce the number of states explored. Since all our 860
 evaluations measured CPU time, they demonstrate that, in many 861
 cases, hybridized search can outperform deeper search using a 862
 single algorithm. 863

B. Systems With New Functionality 864

Although quantitative evaluations help measure and pre- 865
 cisely characterize specific advances that are enabled, our ul- 866
 time goal in this paper is to enable intelligent systems with 867
 functionality that formerly had not been straightforward to 868
 create or which solved problems that could not be solved (in 869
 theory, not just efficiently) in other frameworks. 870

The principal way to demonstrate that Polyscheme can en- 871
 able this is to actually build such systems. The robot from 872
 the last section is a prime example. It can react to changes 873
 while making inferences and finding plans that purely reactive 874
 systems cannot and that conventional inference and reasoning 875
 algorithms achieve only (as we described in the SAT example 876
 given earlier) on very small problem scales. 877

A heterogeneous database retrieval system [25] implemented 878
 in Polyscheme also illustrates its benefits. This system makes 879
 sound and complete inferences over heterogeneous sources of 880
 computation and information. It implements resolution the- 881
 orem proving using common functions, such as forward in- 882
 ference, subgoaling, and identity. Each of these operations is 883
 implemented in specialists based on representations, such as 884
 neural networks, production rules, geospatial coordinates, and 885
 relational databases. If these modules meet certain conditions 886
 (which, in practice, are easy to confirm), the total system’s 887
 inference is sound and complete. This system demonstrates 888
 how hybridizing algorithms provides both the benefits of logic 889
 programming (provable soundness and completeness) and the 890
 efficiency of special representations (such as those from neural 891
 networks and relational databases). 892

Finally, the GenSAT language and GenDPLL algorithm we 893
 mentioned in the motion planning section overcome some 894
 severe difficulties that arise in domains with unknown objects. 895
 Languages, such as GenSAT, that license the inference of 896
 objects that are unknown before inference can lead to models 897
 with infinite numbers of objects. For example, a constraint to 898
 the effect that “all mammals have a mother” and “a person 899
 cannot be their own ancestor” requires a model with infinite 900
 numbers of ancestors for any particular mammal. As another 901
 example of finite theories with infinite models, many context- 902
 free grammars license infinite numbers of derivations. Be- 903
 cause traditional SAT solvers require all objects to be known 904
 in advance—this is true even of lazy SAT solvers such as 905
 LazySAT—they cannot solve many problems in such domains. 906
 In [17], we prove that hybridizing rule matching with weighted 907
 constraint solving in GenDPLL enables models to be found 908
 of GenSAT theories even in many conditions where there are 909
 infinitely many models with infinite numbers of unknown ob- 910
 jects. In [26], we show how to use GenDPLL to provide parses 911
 for probabilistic context-free grammars (PCFGs), even in cases 912
 where a grammar has infinite numbers of derivations. Since 913
 GenSAT can encode an extremely wide variety of constraints, 914
 this result enables linguistic knowledge (in the form of PCFG 915

916 constraints) and nonlinguistic knowledge to jointly constrain
 917 interpretation during parsing. The absence of such an ability
 918 has been a serious obstacle to the use of context to disambiguate
 919 language.

920 The heterogeneous database retrieval and GenSAT results
 921 had been beyond the theoretical reach of existing approaches.
 922 No approach had previously enabled sound and complete an-
 923 swers to queries over information in such a wide variety of
 924 formalisms, and no other approach had enabled reasoning over
 925 constraints as general and flexible as those in SAT to be jointly
 926 reasoned over with grammars that had infinite derivations.
 927 These results thus demonstrate that executing hybrid algorithm
 928 through a focus of attention in Polyscheme not only can speed
 929 inference but also enable inference in situations that had hereto-
 930 fore been theoretically intractable.

931 VII. OTHER APPROACHES TO INTEGRATION

932 Most approaches to integrating algorithms and/or their char-
 933 acteristics into a single system have taken a reductive, mod-
 934 ular, or “fixed” hybrid approach. Reductive approaches tend
 935 to implement an algorithm or solve problems by a reduction
 936 to another approach. This often consigns such systems to the
 937 limitations of the computational formalism or method being
 938 reduced to. For example, reducing a first-order probabilistic
 939 logic reasoning problem to a graphical model belief propa-
 940 gation problem [27] means that one is still limited by the
 941 propositional representation and scalability characteristics of
 942 belief propagation inference algorithms. Modular architectures
 943 for integration tend to enable modules based on different data
 944 structures and algorithms to communicate and/or cooperate.
 945 Normally, the only way to add an algorithm to such a system is
 946 to add a module based on that algorithm. Polyscheme enables
 947 this kind of integration, but it also enables algorithms to be
 948 executed through the focus of attention. This allows every
 949 single step of every algorithm to be executed using many
 950 data structures and algorithms and thus enables a much more
 951 thorough integration of algorithms. Fixed hybrid approaches,
 952 such as Clarion [28] and ACT-R [29], [30], create hybrids
 953 between a few specific algorithms. Both include production
 954 systems that use reinforcement-learning mechanisms for con-
 955 flict resolution but do not enable such close interaction between
 956 other algorithms implemented in those systems.

957 Polyscheme differs from many cognitive architectures in be-
 958 ing primarily inferential and not procedural. Most architectures,
 959 e.g., Icarus [31], Soar [11], Epic [32], and ACT-R, choose an
 960 action at every step. These actions are actual physical actions
 961 or operations on data structures in memory. In Polyscheme, at
 962 every time step, specialists do not take or propose actions but in-
 963 stead take stances on the truth value of propositions and suggest
 964 propositions to attend to. A consequence is that Polyscheme
 965 includes mechanisms for communication and sharing informa-
 966 tion about the truth of propositions that make it much easier to
 967 implement many reasoning and inference algorithms.

968 A focus of attention is a very important part of ACT-R
 969 and Rao’s work [33] on visual routines. However, in ACT-
 970 R, the focus of attention is not used to implement algorithms
 971 and, in neither case, is it feasible (because of the absence of

mechanisms involving truth values and alternate worlds) to 972
 implement and integrate reasoning and inference algorithms 973
 using this sort of focus of attention mechanism. 974

Self-adaptive system architectures (e.g., [34] and [35]) share 975
 some features of this approach. They often contain multiple 976
 algorithms based on different methods. Furthermore, these 977
 architectures often contain some form of self-monitoring that 978
 detects problems among system components and reacts accord- 979
 ingly. This is in some ways similar to the role of choosing focus 980
 to deal with metacognitive problems in Polyscheme. However, 981
 these systems typically do not contain a notion of a focus of 982
 attention that all modules process simultaneously, and they do 983
 not require all components to be able to simulate alternate 984
 worlds. Thus, while nothing prevents these systems from in- 985
 cluding components that implement inference algorithms inside 986
 the modules, they cannot implement these algorithms as the 987
 result of a guided focus of attention that involves all modules 988
 in every step of inference. 989

Polyscheme’s focus of attention is superficially reminiscent 990
 of blackboard systems (e.g., [36]) insofar as it is a shared-entity 991
 multiple-module access. However, the two have numerous and 992
 profound differences: The focus of attention in Polyscheme is 993
 tiny; Polyscheme has no shared memory among modules while 994
 blackboards are such shared memory; Polyscheme modules can 995
 have arbitrarily large, varied, and persistent memories while, in 996
 many blackboard systems, the blackboard is the main form of 997
 memory; blackboard systems have no straightforward way of 998
 exploring alternate states of the world; and, unlike the focus 999
 of attention in Polyscheme, modules in blackboard systems 1000
 generally do not synchronize their operation. Most importantly, 1001
 the lack of shared memory frees modules in Polyscheme to use 1002
 a much more diverse array of representational formalisms. 1003

Finally, like the current approach, Beal [37] proposes that 1004
 many solutions to human intelligence arise from the inter- 1005
 action of specialized processors and explore the value of a 1006
 focus of attention [38]. While it does not implement inference 1007
 algorithms using a focus of attention and simulated worlds, it 1008
 does provide a means of learning interactions among modules. 1009
 In Polyscheme, at present, these interactions are established 1010
 by the system designers, although it is likely that additional 1011
 benefits would result by using such methods to influence the 1012
 interaction of specialists in Polyscheme. 1013

VIII. CONCLUSION

1014

The goal of this paper has been a framework for creating 1015
 a single system that can exhibit the best characteristics of 1016
 algorithms based on different computational formalisms. Our 1017
 approach has been to create a cognitive architecture called 1018
 Polyscheme that can execute hybrids of these algorithms. 1019
 Polyscheme enables two kinds of hybrids. First, a Polyscheme 1020
 system can include modules based on arbitrarily different al- 1021
 gorithms and data structures as long as they implement the 1022
 common functions. Second, and most originally, Polyscheme 1023
 can execute algorithms by guiding the “focus of attention” of 1024
 these modules. Systems created using Polyscheme demonstrate 1025
 that integration through a focus of attention enables systems 1026
 that can make inferences and solve problems in situations 1027
 beyond the reach of existing individual computational methods. 1028

REFERENCES

- 1029
- 1030 [1] M. L. Minsky, "A framework for representing knowledge," in *The Psychology of Computer Vision*, P. H. Winston, Ed. New York: McGraw-Hill, 1975.
- 1031
- 1032
- 1033 [2] R. C. Schank and R. P. Abelson, *Scripts, Plans, Goals, and Understanding: An Inquiry Into Human Knowledge Structures*. Hillsdale, NJ: Lawrence Erlbaum, 1977.
- 1034
- 1035
- 1036 [3] J. Kolodner, *Case-Based Reasoning*. Menlo Park, CA: Morgan Kaufman, 1993.
- 1037
- 1038 [4] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-6, no. 6, pp. 721–741, Nov. 1984.
- 1039
- 1040
- 1041 [5] B. Selman, H. Levesque, and D. Mitchell, "A new method for solving hard satisfiability problems," in *Proc. 10th Nat. Conf. Artif. Intell.*, 1992, pp. 440–446.
- 1042
- 1043
- AQ8 1044 [6] N. Een and N. Sorensson, "MiniSat—A SAT solver with conflict-clause minimization," in *Proc. SAT Competition*, 2005.
- 1045
- 1046 [7] M. Moskewicz, C. Madigan, Y. Zhao, L. Zhang, and S. Malik, "Chaff: Engineering an efficient SAT solver," in *Proc. 39th Des. Autom. Conf.*, Las Vegas, NV, 2001, pp. 530–535.
- 1047
- 1048
- 1049 [8] M. Davis, G. Logemann, and D. Loveland, "A machine program for theorem proving," *Commun. ACM*, vol. 5, no. 7, pp. 394–397, Jul. 1962.
- 1050
- 1051 [9] J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proc. Nat. Acad. Sci.*, vol. 79, no. 8, pp. 2554–2558, Apr. 1982.
- 1052
- 1053
- 1054 [10] M. I. Jordan and C. Bishop, "Neural networks," in *CRC Handbook of Computer Science*, A. B. Tucker, Ed. Boca Raton, FL: CRC Press, 1997.
- 1055
- 1056 [11] J. E. Laird, A. Newell, and P. S. Rosenbloom, "Soar: An architecture for general intelligence," *Artif. Intell.*, vol. 33, no. 1, pp. 1–64, Sep. 1987.
- 1057
- 1058 [12] N. L. Cassimatis, "Reasoning as cognitive self-regulation," in *Integrated Models of Cognitive Systems*, W. Gray, Ed. New York: Oxford Univ. Press, 2007.
- 1059
- 1060
- AQ9 1061 [13] A. M. N. Cassimatis, "Reasoning as perceptual simulation," in *Cognitive Processing*, to be published.
- 1062
- 1063 [14] N. L. Cassimatis, J. G. Trafton, M. Bugajska, and A. C. Schultz, "Integrating cognition, perception, and action through mental simulation in robots," *Robot. Auton. Syst.*, vol. 49, no. 1/2, pp. 13–23, Nov. 2004.
- 1064
- 1065
- 1066 [15] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Trans. Syst. Sci. Cybern.*, vol. SSC-4, no. 2, pp. 100–107, Jul. 1968.
- 1067
- 1068
- 1069 [16] S. L. Epstein, "Pragmatic navigation: Reactivity, heuristics, and search," *Artif. Intell.*, vol. 100, no. 1/2, pp. 275–322, Apr. 1998.
- 1070
- 1071 [17] A. M. N. Cassimatis and P. Bignoli, "Inference with relational theories over infinite domains," in *Proc. FLAIRS*, 2009.
- 1072
- AQ10 1073 [18] P. Singla and P. Domingos, "Memory-efficient inference in relational domains," in *Proc. AAAI*, 2006, pp. 488–493.
- 1074
- 1075 [19] F. Heras, J. Larrosa, and A. Oliveras, "MiniMaxSAT: An efficient weighted Max-SAT solve," *J. Artif. Intell. Res.*, vol. 31, pp. 1–32, 2008.
- 1076
- 1077 [20] S. Wintermute and J. E. Laird, "Predicate projection in a bimodal spatial reasoning system," in *Proc. 22nd AAAI Conf. Artif. Intell.*, Vancouver, BC, Canada, 2007, pp. 1572–1577.
- 1078
- 1079
- 1080 [21] D. H. Wolpert and W. G. Macready, *No Free Lunch Theorems for Search*. Santa Fe, NM: Santa Fe Inst., 1995.
- 1081
- 1082 [22] D. H. Wolpert and W. G. Macready, "Coevolutionary free lunches," *IEEE Trans. Evol. Comput.*, vol. 9, no. 6, pp. 721–735, Dec. 2005.
- 1083
- 1084 [23] J. Gu, "Efficient local search for very large-scale satisfiability problems," *SIGART Bull.*, vol. 3, no. 1, pp. 8–12, Jan. 1992.
- 1085
- 1086 [24] N. L. Cassimatis, "A cognitive substrate for human-level intelligence," *Artif. Intell. Mag.*, vol. 27, no. 2, pp. 45–56, Jul. 2006.
- 1087
- 1088 [25] N. L. Cassimatis, "A framework for answering queries using multiple representation and inference technique," in *Proc. 10th Int. Workshop Knowl. Representation Meets Databases*, 2003.
- 1089
- AQ11 1090 [26] A. Murugesan, N. L. Cassimatis, S. Dugas, and M. Bugajska, "Combining probabilistic context-free parsing with general inference," in *Proc. UAI*, Vancouver, BC, Canada, 2007.
- 1091
- 1092
- 1093
- [27] P. Domingos and M. Richardson, "Markov logic networks," *Mach. Learn.*, vol. 62, no. 1/2, pp. 107–136, Feb. 2006.
- [28] R. Sun, "The CLARION cognitive architecture: Extending cognitive modeling to social simulation," in *Cognition and Multi-Agent Interaction*. New York: Cambridge Univ. Press, 2004.
- [29] J. R. Anderson, "Human symbol manipulation within an integrated cognitive architecture," *Cogn. Sci.*, vol. 29, no. 3, pp. 313–341, May 2005.
- [30] J. R. Anderson and C. Lebiere, *The Atomic Components of Thought*. Hillsdale, NJ: Lawrence Erlbaum Associates, 1998.
- [31] P. Langley and D. Choi, "A unified cognitive architecture for physical agents," in *Proc. 21st Nat. Conf. Artif. Intell.*, Boston, MA, 2006, pp. 1469–1474.
- [32] D. Kieras and D. E. Meyer, "An overview of the EPIC architecture for cognition and performance with application to human-computer interaction," *Hum.-Comput. Interact.*, vol. 12, no. 4, pp. 291–438, Dec. 1997.
- [33] S. Rao, "Visual routines and attention," in *Electrical Engineering and Computer Science*. Cambridge, MA: MIT Press, 1998, p. 86.
- [34] H. E. Shrobe, R. Laddaga, R. Balzer, N. M. Goldman, D. Wile, M. Tallis, T. Hollebeek, and A. Egyed, "Self-adaptive systems for information survivability: PMOP and AWD RAT," in *Proc. SASO*, 2007, pp. 332–335.
- [35] R. Laddaga, P. Robertson, and H. E. Shrobe, "Introduction to self-adaptive software: Applications," in *Proc. IWSAS*, 2001, pp. 1–5.
- [36] B. Hayes-Roth, "A blackboard architecture for control," *Artif. Intell.*, vol. 26, no. 3, pp. 251–321, Jul. 1985.
- [37] J. Beal, "Learning by learning to communicate," in *Electrical Engineering and Computer Science*. Cambridge, MA: MIT Press, 2007.
- [38] J. Beal, "Shared focus of attention for heterogeneous agents," in *Proc. 7th Int. Conf. AAMAS*, 2008, pp. 1627–1630.
- Nicholas Cassimatis, photograph and biography not available at the time of publication. 1123 AQ12 1124
- Perrin Bignoli, photograph and biography not available at the time of publication. 1125 AQ13 1126
- Magdalena Bugajska, photograph and biography not available at the time of publication. 1127 AQ14 1128
- Scott Dugas, photograph and biography not available at the time of publication. 1129 AQ15
- Unmesh Kurup, photograph and biography not available at the time of publication. 1130 AQ16 1131
- Arthi Murugesan, photograph and biography not available at the time of publication. 1132 AQ17 1133
- Paul Bello, photograph and biography not available at the time of publication. 1134 AQ18

AUTHOR QUERIES

AUTHOR PLEASE ANSWER ALL QUERIES

AQ1 = Please validate the address and postal code provided for “Rensselaer Polytechnic Institute.”

AQ2 = Please validate the address and postal code provided for “Office of Naval Research.”

AQ3 = “Davis-Putnam-Logemann-Loveland” was taken as the expanded form of the acronym “DPLL.” Please check if appropriate.

AQ4 = “Case-based reasoning” was taken as the expanded form of the acronym “CBR.” Please check if correct.

AQ5 = “Rule specialist infers that B...” was changed to “rule specialist infers that B is true.” Please check if appropriate.

AQ6 = The acronym “NP” was defined as “nondeterministic polynomial time.” Please check if correct.

AQ7 = “Probabilistic context-free grammar” was taken as the expanded form of the acronym “PCFG.” Please check if appropriate.

AQ8 = Please provide page range in Ref. [6].

AQ9 = Please provide publication update in Ref. [13].

AQ10 = Please provide page range in Ref. [17].

AQ11 = Please provide page range in Ref. [25].

AQ12 = Please provide photograph and biography of the author “Nicholas Cassimatis.”

AQ13 = Please provide photograph and biography of the author “Perrin Bignoli.”

AQ14 = Please provide photograph and biography of the author “Magdalena Bugjaska.”

AQ15 = Please provide photograph and biography of the author “Scott Dugas.”

AQ16 = Please provide photograph and biography of the author “Unmesh Kurup.”

AQ17 = Please provide photograph and biography of the author “Arthi Murugesan.”

AQ18 = Please provide photograph and biography of the author “Paul Bello.”

END OF ALL QUERIES