

Middlesex University Research Repository

An open access repository of

Middlesex University research

<http://eprints.mdx.ac.uk>

Windridge, David ORCID: <https://orcid.org/0000-0001-5507-8516> (2017) Emergent intentionality in perception-action subsumption hierarchies. *Frontiers in Robotics and AI*, 4 . ISSN 2296-9144 (doi:10.3389/frobt.2017.00038)

Final accepted version (with author's formatting)

This version is available at: <http://eprints.mdx.ac.uk/22394/>

Copyright:

Middlesex University Research Repository makes the University's research available electronically.

Copyright and moral rights to this work are retained by the author and/or other copyright owners unless otherwise stated. The work is supplied on the understanding that any use for commercial gain is strictly forbidden. A copy may be downloaded for personal, non-commercial, research or study without prior permission and without charge.

Works, including theses and research projects, may not be reproduced in any format or medium, or extensive quotations taken from them, or their content changed in any way, without first obtaining permission in writing from the copyright holder(s). They may not be sold or exploited commercially in any format or medium without the prior written permission of the copyright holder(s).

Full bibliographic details must be given when referring to, or quoting from full items including the author's name, the title of the work, publication details where relevant (place, publisher, date), pagination, and for theses or dissertations the awarding institution, the degree type awarded, and the date of the award.

If you believe that any material held in the repository infringes copyright law, please contact the Repository Team at Middlesex University via the following email address:

eprints@mdx.ac.uk

The item will be removed from the repository while any claim is being investigated.

See also repository copyright: re-use policy: <http://eprints.mdx.ac.uk/policies.html#copy>

Emergent Intentionality in Perception-Action Subsumption Hierarchies

David Windridge^{1,*}

¹Department of Computer Science, Middlesex University, London, UK

Correspondence*:

David Windridge

d.windridge@mdx.ac.uk

2 ABSTRACT

3 A *cognitively-autonomous* artificial agent may be defined as one able to modify both its external
4 world-model *and* the framework by which it represents the world, requiring two simultaneous
5 optimization objectives. This presents deep epistemological issues centered on the question of
6 how a framework for representation (as opposed to the entities it represents) may be objectively
7 validated. In this summary paper, formalizing previous work in this field, it is argued that sub-
8 sumptive perception-action learning has the capacity to resolve these issues by *a)* building the
9 perceptual hierarchy from the bottom up so as to ground all proposed representations and *b)*
10 maintaining a bijective coupling between proposed percepts and projected action possibilities to
11 ensure empirical falsifiability of these grounded representations. In doing so, we will show that
12 such subsumptive perception-action learners intrinsically incorporate a model for how intentiona-
13 lity emerges from randomized exploratory activity in the form of 'motor babbling'. Moreover, such
14 a model of intentionality also naturally translates into a model for human-computer interfacing
15 that makes minimal assumptions as to cognitive states.

16 **Keywords:** Perception-Action Learning, Intention recognition, Embodied cognition, Subsumption Hierarchies, Symbol Grounding

1 INTRODUCTION

17 Significant deficits have been apparent in traditional approaches to embodied computer vision for some
18 time Dreyfus (1972). In the conventional approach to autonomous robotics, a computer vision system is
19 employed to build a model of the agent's environment *prior* to the act of planning the agent's actions within
20 the modeled domain. Visuo-haptic data arising from these actions will then typically be used to further
21 constrain the environment model, either actively or passively (in active learning the agent actions are driven
22 by the imperative of reducing ambiguity in the environment model Settles (2010); Koltchinskii (2010)).

23 However, it is apparent, in this approach, that there exists a very wide disparity between the visual
24 parameterization of the agent's domain and its action capabilities within it Nehaniv et al. (2002). For
25 instance, the agent's visual parametric freedom will typically encompass the full intensity ranges of the
26 RGB channels of each individual pixel of a camera CCD, such the the range of *possible* images generated
27 per time-frame is of an extremely large order of magnitude, despite the fact that only a minuscule fraction
28 of this representational space would ever be experienced by the agent. (Note that this observation is not
29 limited purely to vision based approaches - alternative modalities such as LIDAR and SONAR would also
30 exhibit the same issues). On the other hand, the agent's motor capability is likely to be very much more

31 parametrically-constrained (perhaps consisting of the possible Euler angle settings of the various actuator
 32 motors). This disparity is manifested in classical problems such as *framing* McCarthy and Hayes (1969)
 33 and *symbol grounding*. (The latter occurs when abstractly-manipulated symbolic objects lack an intrinsic
 34 connection to the real-world objects that they represent; thus a chess-playing robot typically requires a
 35 prior supervised computer vision to be solved in order to apply deduced moves to visually-presented chess
 36 pieces.)

37 Perception-Action (P-A) learning was proposed in order to overcome these issues, adopting as its informal
 38 motto, ‘action *precedes* perception’ Granlund (2003); Felsberg et al. (2009). By this it is meant that, in a
 39 fully-formalizable sense, actions are conceptually prior to perceptions; i.e. perceptual capabilities should
 40 depend on action-capabilities and not vice versa. (We thus distinguish *PA-learning* from more generalized
 41 forms of learning within a perception/action context (cf. e.g. d. R. Millan (2016); Mai et al. (2013); Masuta
 42 et al. (2015)), in which the nature of the perceptual domain remains fixed *a priori* [albeit with potential
 43 variations in e.g. visual saliency]).

44 It will be the argument of this paper that perception-action learning, as well as having this capacity to
 45 resolve fundamental epistemic questions about emergent representational capacity, also naturally gives
 46 a model for emergent intentionality that applies to both human and artificial agents, and may thus be
 47 deployed as an effective design-strategy in human-computer interfacing.

2 PERCEPTION-ACTION LEARNING

48 Perception-Action learning agents thus proceed by randomly sampling their action space (‘motor babbling’).
 49 For each motor action that produces a discernible perceptual output in the bootstrap representation space
 50 S (consisting of e.g. camera pixels), a percept $p_i \in S$ is greedily allocated. The agent thus progressively
 51 arrives at a set of novel percepts that relate directly to the agent’s action capabilities in relation to the
 52 constraints of the environment (i.e. the environment’s *affordances*); the agent learns to perceive only
 53 that which it can change. More accurately, the agent learns to perceive only that which it *hypothesizes*
 54 that it can change - thus, the set of experimental data points $\cup_i p_i \subset S$ can, in theory, be generalized
 55 over so as to create an *affordance-manifold* that can be mapped onto the action space via the injective
 56 relation $\{actions\} \rightarrow \{percept_{initial}\} \times \{percept_{final}\}$ Windridge and Kittler (2010, 2008); Windridge
 57 et al. (2013a).

58 2.1 Subsumptive Perception-Action Learning

59 Importantly, this approach permits *Cognitive Bootstrapping* Windridge and Kittler (2010), the boot-
 60 strapping of an autonomous agent’s representational framework simultaneously with the world-model
 61 represented in terms of that framework. This centers on the fact that the learned manifold embodying
 62 the *injective* relation $\{actions\} \rightarrow \{percept_{initial}\} \times \{percept_{final}\}$ represents a constrained subset of the
 63 initial action domain, and as such, is susceptible to parametric compression. Furthermore, this parametric
 64 compression in the action domain (corresponding to the bootstrapping of a higher level action) necessarily
 65 corresponds to a parametric compression in the perceptual domain (P-A learning enforces a *bijective*
 66 relation $\{actions\} \rightarrow \{percept_{initial}^{new}\} \times \{percept_{final}^{new}\}$ such that each hypothesizable action (ie intention
 67 primitive) has a unique, discriminable outcome Windridge and Kittler (2010, 2008); Windridge et al.
 68 (2013a)).

69 Each induced higher-level action/intention (e.g. *Translate*) is thus created co-extantly with a higher-
 70 level percept domain (e.g. *Object*). The falsifiability of such induced representational concepts arises

71 from actively addressing the question of whether this higher-level perception in fact constitutes a useful
72 description of the world i.e. whether it yields a net compression in the agent's internal representation of
73 its own possible interactions with the world (its affordances). In particular, it is argued in Windridge and
74 Kittler (2008), that the perception-action bijectivity constraint applied in such a hierarchical manner is
75 uniquely sufficient to enable simultaneous empirical falsifiability of the cognitive agent's world model
76 and the means by which this world is perceived (by virtue of the implicit grounding of the unique set of
77 higher-level percepts so generated).

78 Very often parametric compressibility will be predicated on the discovery of *invariances* in the existing
79 perceptual space with respect to randomized exploratory actions. Thus, for example, an agent might
80 progress from a pixel-based representation of the world to an object-based representation of the world
81 via the discovery that certain patches of pixels *retain their (relative) identity* under translation, i.e. such
82 that it becomes far more efficient to represent the world in terms of indexed objects rather than pixel
83 intensities (though the latter would, of course, still constitute the base of the representational hierarchy).
84 This particular representational enhancement can represent an enormous compression Wolff (1987); a pixel-
85 based representation has a parametric magnitude of P^n (with P and n being the intensity resolution and
86 number of pixels, respectively), while an object-based representation typically has a parametric magnitude
87 of $\sim n^o$, $o \ll n$, where o is the number of objects.

88 When such a high level perceptual manifold is created it permits proactive sampling - the agent can
89 propose actions with perceptual outcomes that have not yet been experienced by the agent, but which are
90 consistent with its current representational model (this guarantees falsifiability of both the perceptual model
91 as well as the generalized affordance model). Perception-Action learning thus constitutes a form of active
92 learning: randomized selection of perceptual goals within the hypothesized perception-action manifold
93 leads more rapidly to the capture of data that might falsify the current hypothesis than would otherwise be
94 the case (i.e. if the agent were performing randomly-selected actions within in the original motor domain).
95 Thus, while the system is always 'motor babbling' in a manner analogous to the learning process of infant
96 humans, the fact of carrying out this motor babbling in a higher-level P-A manifold means that the learning
97 system as a whole more rapidly converges on the "correct" model of the world. (Correct in the sense of
98 being a true model of the world's affordances; i.e. every possible instantiation of the induced high-level
99 actions terminates in the anticipated percept, with no possible environmental actions being overlooked.)

100 This P-A motor-babbling activity can take place in *any* P-A manifold, of whatever level of abstraction;
101 we may thus, by combining the idea of P-A learning with Brooke's notion of task subsumption, conceive
102 of a *hierarchical Perception-Action learner* (Shevchenko et al. (2009)), in which a vertical representation
103 hierarchy is progressively constructed for which randomized exploratory motor activity at the highest level
104 of the corresponding motor hierarchy would rapidly converge on an ideal representation of the agent's world
105 in terms of its affordance potentialities. Such a system would thus converge upon both a model of the world,
106 and an ideal strategy for representation of that world in terms of the learning agent's action capabilities
107 within it. In the example given, which juxtaposes a simulated camera-equipped robot arm in relation to a
108 child's shape-sorter puzzle, the robotic agent commences by motor babbling in the initial motor-actuator
109 domain, and eventually progresses to motor-babbling in the bootstrapped "move-shape-to-hole" action
110 domain (i.e. placing a randomly chosen object into its corresponding hole). This apparently intentional
111 activity amounts to solving the shape-sorter puzzle, even though the system is still only motor babbling

112 albeit at a higher level of the induced hierarchy, and has no prior programming as to the 'goal' of the
113 environment¹.

114 Procedurally, this takes place as a recursive loop alternating between *exploration*, *generalization* and
115 *representation* as in Algorithm 1. Note in particular, in Algorithm 1, that the act of parametrically-
116 instantiating a proposed bijective perception-action term $\{P_n^{\text{initial}}\} \times \{P_n^{\text{final}}\}$ with respect to an initial
117 perceptual state $\{P_n^{\text{initial}}\}$ and a sought perceptual end-state $\{P_n^{\text{final}}\}$ is equivalent to formulating an *intention*
118 (which may or may not be achievable in the environment).

Algorithm 1 *Ab Initio* Induction of Perception-Action Hierarchy in Artificial Agents

- 1: **Initialization** Obtain:
 - 2: Bootstrap percept set $\{P_1\}$ (eg camera pixel)
 - 3: Bootstrap action set $\{A_1\}$ (motor primitives)
 - 4: Inference mechanism capable of generalizing exploratory samples from function M
 $M : \{P_n^{\text{initial}}\} \times \{P_n^{\text{final}}\} \times \{A\} \rightarrow \{\text{achieved}, \text{not_achieved}\}$
 - 5: **while** prediction accuracy < threshold) **do**
 - 6: **A) Carry-out randomized exploratory activity on basis of representational-framework**
 i.e. generate grounded top-down parametric instantiations
 $A_{i \leq n}(P^{\text{initial}}, P^{\text{final}})$ by randomly selecting initial &
 target percepts at *proposed* top level of hierarchy, n
 - 7: **B) Induce rules governing action legitimacy**
 legitimate actions achieve intended perceptual goal
 (= affordance-based model of world)
 Generate function $M : \{P\} \times \{P\} \times \{A\} \rightarrow \{\text{true}, \text{false}\}$
 (e.g. via first-order logical induction or stochastic discrimination)
 - 8: **C) Remap perceptual variables to represent novel high-level action hypothesis in most efficient manner**
 i.e. form the bijection: $\{A_{n+1}\} \leftrightarrow \{P_{n+1}\} \times \{P_{n+1}\}$
 - 9: **end while**
-

119 Perceptual goals thus exist at all levels of the hierarchy, and the subsumptive nature of the hierarchy
120 means that goals and sub-goals are scheduled with increasingly specific content as the high-level abstract
121 goal is progressively grounded through the hierarchy². (Thus, as humans, we may conceive the high-level
122 intention 'drive to work', which in order to be enacted, involves the execution of a large range of sub-goals
123 with correspondingly lower-level perceptual goals e.g. the intention 'stay in the center of the lane', etc).
124 (The hierarchical perception-action paradigm at no stage specifies *how* the scheduled sub-task is to achieve

¹ In this case, the "move-shape-to-hole" action is induced following the failure of the "move-shape-onto-surface" action to produce the anticipated result (i.e. when, following exploratory "move-shape-onto-surface" actions, the object happens by chance to fall into a hole to which matches its shape and orientation). This immediately falsifies the existing rule-base such that that the action domain is necessarily modified, by first-order logical rule-induction, to account for this possibility. In the context of the PA bijectivity condition this requires the existence of an action "move-shape-to-hole" perceptually parameterized by a set of labels corresponding to the perceptual representation of holes. Exploratory instantiation of this higher-level rule then corresponds to placing random objects into their corresponding holes i.e. "solving the shapersorter", even though no such external goal specification has taken place.

² The subsumption hierarchy thus acts bidirectionally; the hierarchy is learned bottom-up, while exploratory actions are instantiated top-down.

125 the perceptual goal - this is free within the framework, and may be achieved by a variety of mechanisms
126 e.g. optimal control, minimum jerk etc).

127 Moreover, these perceptual goals have no internal content; in a fully-achieved perception-action learning
128 agent, the environment effectively “becomes its *own* representation”, Newell and Simon (1976), representing
129 a significant compression of the information that an agent needs to retain. This relates directly to the
130 issue of symbol grounding, a seminal problem in the conceptual underpinning of the classical approach
131 to machine learning Harnad (1990). The problem arises when one attempts to relate an abstract symbol
132 manipulation system (it was a common historical assumption that computational reasoning would center
133 on a system such as first-order logic deduction McCarthy and Hayes (1969)) with the stochastic, shifting
134 reality of sensor data. In hierarchical P-A learning the problem is eliminated by virtue of the fact that
135 symbolic representations are *abstracted from the bottom-up* Marr (1982); Gärdenfors (1994); Modayil
136 (2005); Granlund (2003). They are thus always intrinsically grounded (for an example of utilization of
137 first-order logic induction within a subsumption hierarchy see Windridge and Kittler (2010)).

138 The subsumption hierarchy is thus typically characterized by continuous stochastic relationships on the
139 lower levels with more discrete, symbolic manipulations occurring at the higher levels - for this reason,
140 consistent with findings of Shevchenko et al. (2009), motor-babbling at the top of the representation
141 hierarchy involves the spontaneous scheduling of perceptual goals and sub-goals at the lower level of the
142 hierarchy in a way that (as the hierarchy becomes progressively deeper) looks increasingly *intentional*.
143 (This phenomenon is readily apparent in the development of motor movement of human infants as schema
144 abstraction takes place - for instance, the intuition of a generalized percept category *container* correlates
145 with the attempt to validate this notion via the repeated placing of a variety of objects into a variety of
146 containers; cf Hintzman (1986) for an analysis of scheme abstraction in infants).

147 Such high-level schema-employment in humans can, in principle, be detected via an appropriate
148 classification system, enabling novel forms of intentional interfacing between humans and machines.

3 HUMAN-COMPUTER INTERFACING

149 The percept-action relationship may thus be modeled in reverse to characterize human intentional behavior;
150 consider how, as humans we typically represent our environment when driving a vehicle. At one level, we
151 internally represent the immediate environment in metric-related terms (i.e. we are concerned with our
152 proximity to other road users, to the curb and so on). At a higher level, however, we are concerned primarily
153 with *navigation*-related entities (i.e. how individual roads are *connected*). That the latter constitutes a higher
154 hierarchical level, both mathematically and experientially, is guaranteed by the fact that the topological
155 representation *subsumes*, or supervenes upon, the metric representation; i.e. the metric-level provides
156 additional ‘fine-grained’ information to the road topology: the metric representation can be reduced to the
157 topological representation, but not vice versa.

158 We can thus adopt the perception-action bijectivity principle as a *design paradigm* in building HCI
159 systems by demanding that intentional acts on the part of the user are correlated maximally-efficiently
160 (i.e. bijectively) with perceptual transitions apparent to the user. This thus permits a user interface that
161 makes minimal assumptions as to underlying cognitive processes, assuming nothing more than the ability
162 to discriminate percept terminia. This subsumption architecture paradigm was used in Windridge et al.
163 (2013b) to demonstrate, in the context of a driver assistance system, induction of the intentional hierarchy
164 for drivers of a vehicle in which action and eye-gaze take place with respect an external road camera view.
165 The corresponding system constructed for the the project demonstrator was thus able to determine the

166 driver's intentional hierarchy in relation to the current road situation and provide assistance accordingly.
167 In principle, such an interface can also be extended to direct mechanical assistance by substituting the
168 computationally modeled perception-action system for the human perception-action system along the lines
169 of the horse-rider interaction paradigm.

170 Such P-A HCI interfaces will generally require the ability to adaptively link high-level reasoning processes
171 (modeled by e.g. first-order logic) with low-level reactive processes (modeled, for example, stochastically).
172 This amounts to a requirement to propagate learning across the symbolic/sub-symbolic divide. However,
173 because the P-A hierarchy does not make intrinsic distinction between these (there is only progressively
174 grounded P-A abstraction), it is possible to conceive of P-A learning platforms that embody a variety of
175 different learning approaches at different hierarchical levels, but which are all able to learn together by
176 passing derivatives between hierarchical layers in a manner analogous to deep learning approaches.

177 An example utilizing a two-layer P-A hierarchy is given in Windridge et al. (2013a) which incorporates a
178 fuzzy first-order logic reasoning process on the top level and an Euler-Lagrange based trajectory optimisation
179 process on the lower level. The fuzzy-reasoning process employs predicates embodying the P-A bijectivity
180 condition to compute the fixed point of the logical operator T_P ; i.e. $T_P(I) = I$ for each time interval t .

181 I is thus the *Herbrand* model, the minimal logically-consistent 'world model' for time t , of the logical
182 programme P (where $P = \text{fixed clauses} + \text{temporalized detections} + \text{ground atom queries for } t + 1$; P hence
183 embodies a series of first-order logical rules concerning traffic behavior). This functionalization of the
184 logical reasoning enables the predicate-prediction disparity with respect to the lower-level to modulate
185 the lower-level's Euler-Lagrange optimization via the inter-level Jacobean derivatives. The net result is
186 logically-weighted updating of the Euler-Lagrange optimization that allows for on-line (top-down and
187 bottom-up) adaptivity to human inputs. For example, in top-down terms, this allows a logically-influenced
188 Bayesian prior for gaze-location at junctions to be derived. It also allows for adaptive symbol tethering; for
189 example actively associating eye-gaze clusters with specific semantically-described road entities (such as
190 stop & give-way signs) via their logical context.

191 In principle, any high-level abstract reasoning or induction process can be incorporated with low-level
192 stochastic learning in this manner; highly flexible human-computer interfaces are thus made possible
193 through adopting perception-action bijectivity as a *design principle*.

4 CONCLUSION

194 We have proposed perception-action hierarchies as a natural solution to the problem of representational
195 induction in artificial agents in a manner that maintains empirical validatability. In such *ab initio* P-A
196 hierarchies (i.e. where cognitive representations are bootstrapped in a bottom-up fashion), exploration is
197 conducted via motor-babbling at progressively higher levels of the hierarchy. This necessarily involves the
198 spontaneous scheduling of perceptual goals and sub-goals in the induced lower levels of the hierarchy in
199 such a way that, as the hierarchy becomes deeper, that the randomized exploration becomes increasingly
200 'intentional' (a phenomenon that is readily apparent in the development of motor movement in human
201 infants).

202 This has implications for social robotics; in particular, it becomes possible to envisage communicable
203 actions within collections of agents employing P-A hierarchies. Here, the same bijectivity considerations
204 apply to perceptions and actions as before, however the induction and grounding of symbols would be
205 conducted through linguistic exchange (we note in passing that the perception-action bijectivity constraint

206 implicitly embodies the notion of mirroring without requiring specific perceptual apparatus - ‘mirror neurons’
207 etc).

208 P-A subsumption hierarchies naturally also encompass symbolic/sub-symbolic integration and permit
209 adaptive learning with respect to existing knowledge bases; in this case a bijective P-A consistency criterion
210 is imposed on the engineered subsumption hierarchy. Moreover, P-A-subsumption hierarchies naturally
211 lend themselves to a “deep” formulation in neural-symbolic terms d’Avila Garcez et al. (2009); this is the
212 subject of ongoing research.

213 We therefore conclude that Perception-Action learning, as well as enabling autonomous cognitive
214 bootstrapping architectures, also constitutes a particularly straightforward approach to modeling human
215 intentionality, in that it makes fewest cognitive assumptions - the existence of perceptual representation is
216 only assumed in so far as it directly relates to an observable high-level action concept (such a ‘navigating a
217 junction’, ‘stopping at a red light’, etc); conversely, the ability to correctly interpret a human agent’s action
218 implicitly invokes a necessary and sufficient set of perceptual representations on the part of the agent. This
219 bijectivity of perception and action also gives a natural explanation for wider intention-related phenomenon
220 such as *action mirroring*.

CONFLICT OF INTEREST STATEMENT

221 The authors declare that the research was conducted in the absence of any commercial or financial
222 relationships that could be construed as a potential conflict of interest.

FUNDING

223 The first author would like to acknowledge financial support from the Horizon 2020 European Research
224 project DREAMS4CARS (no. 731593).

REFERENCES

- 225 d. R. Millan, J. (2016). Brain-controlled devices: the perception-action closed loop. In *2016 4th*
226 *International Winter Conference on Brain-Computer Interface (BCI)*. 1–2. doi:10.1109/IWW-BCI.
227 2016.7457451
- 228 d’Avila Garcez, A. S., Lamb, L. C., and Gabbay, D. M. (2009). *Neural-Symbolic Cognitive Reasoning*.
229 Cognitive Technologies (Springer). doi:10.1007/978-3-540-73246-4
- 230 Dreyfus, H. (1972). *What Computers Can’t Do* (New York: Harper and Row)
- 231 Felsberg, M., Wiklund, J., and Granlund, G. (2009). Exploratory learning structures in artificial cognitive
232 systems. *Image and Vision Computing* 27, 1671–1687
- 233 Gärdenfors, P. (1994). How logic emerges from the dynamics of information. *Logic and Information Flow*
234 , 49–77
- 235 Granlund, G. (2003). Organization of architectures for cognitive vision systems. In *Proceedings of*
236 *Workshop on Cognitive Vision* (Schloss Dagstuhl, Germany)
- 237 Harnad, S. (1990). The symbol grounding problem. *Physica D* , 335–346
- 238 Hintzman, D. L. (1986). Schema abstraction in a multiple-trace memory model. *Psychological review* 93,
239 411–428
- 240 Koltchinskii, V. (2010). Rademacher complexities and bounding the excess risk in active learning. *The*
241 *Journal of Machine Learning Research* 11, 2457–2485

- 242 Mai, X., Zhang, X., Jin, Y., Yang, Y., and Zhang, J. (2013). Simple perception-action strategy based on
243 hierarchical temporal memory. In *2013 IEEE International Conference on Robotics and Biomimetics*
244 (*ROBIO*). 1759–1764. doi:10.1109/ROBIO.2013.6739722
- 245 Marr, D. (1982). *Vision: A Computational Approach* (San Fr.: Freeman & Co.)
- 246 Masuta, H., o. Lim, H., Motoyoshi, T., Koyanagi, K., and Oshima, T. (2015). Direct perception and action
247 system for unknown object grasping. In *2015 24th IEEE International Symposium on Robot and Human*
248 *Interactive Communication (RO-MAN)*. 313–318. doi:10.1109/ROMAN.2015.7333637
- 249 McCarthy, J. and Hayes, P. (1969). Some philosophical problems from the standpoint of artificial
250 intelligence. *Machine Intelligence* , 463–502
- 251 Modayil, J. (2005). Bootstrap learning a perceptually grounded object ontology. Retr. 9/5/2005
252 <http://www.cs.utexas.edu/users/modayil/modayil-proposal.pdf>
- 253 Nehaniv, C. L., Polani, D., Dautenhahn, K., te Boekhorst, R., and Canamero, L. (2002). Meaningful
254 information, sensor evolution, and the temporal horizon of embodied organisms. In *Artificial Life VIII*,
255 ed. B. Standish, Abbass (MIT Press). 345–349
- 256 Newell, A. and Simon, H. (1976). The theory of human problem solving; reprinted in collins & smith
257 (eds.). In *Readings in Cognitive Science, section 1.3*.
- 258 Settles, B. (2010). Active learning literature survey. *University of Wisconsin, Madison*
- 259 Shevchenko, M., Windridge, D., and Kittler, J. (2009). A linear-complexity reparameterisation strategy for
260 the hierarchical bootstrapping of capabilities within perception–action architectures. *Image and Vision*
261 *Computing* 27, 1702–1714
- 262 Windridge, D., Felsberg, M., and Shaukat, A. (2013a). A framework for hierarchical perception-action
263 learning utilizing fuzzy reasoning. *Cybernetics, IEEE Transactions on* 43, 155–169. doi:10.1109/
264 TSMCB.2012.2202109
- 265 Windridge, D. and Kittler, J. (2008). Epistemic constraints on autonomous symbolic representation
266 in natural and artificial agents. In *Studies in Computational Intelligence: Applications of Com-*
267 *putational Intelligence in Biology* (Springer Berlin Heidelberg), vol. 122. 395–422. doi:10.1007/
268 978-3-540-78534-7_16
- 269 Windridge, D. and Kittler, J. (2010). Perception-action learning as an epistemologically-consistent model
270 for self-updating cognitive representation. In *Brain Inspired Cognitive Systems 2008* (Springer). 95–134
- 271 Windridge, D., Shaukat, A., and Hollnagel, E. (2013b). Characterizing driver intention via hierarchical
272 perception-action modeling. *Human-Machine Systems, IEEE Transactions on* 43, 17–31. doi:10.1109/
273 TSMCA.2012.2216868
- 274 Wolff, J. G. (1987). Cognitive development as optimisation. In *Computational Models of Learning*, ed.
275 L. Bolc (Heidelberg: Springer-Verlag), 161–205