
Table of Contents

| | |
|---|-----------|
| Chapter 1: Introduction | 15 |
| 1.1 What is an adaptive autonomous agent? | 15 |
| 1.2 Some basic concepts | 17 |
| 1.3 Focus of attention in learning | 18 |
| Chapter 2: The Basic Framework | 23 |
| 2.1 Focus of Attention Methods | 23 |
| 2.1.1 Introduction | 23 |
| 2.1.2 Perceptual selectivity | 25 |
| 2.1.3 Cognitive selectivity | 26 |
| 2.2 The testbed microworlds | 28 |
| 2.2.1 The infant/eyehand scenario | 29 |
| 2.2.2 The Hamsterdam scenario | 32 |
| Chapter 3: Goal-independent Learning | 39 |
| 3.1 Model of the learning system | 39 |
| 3.1.1 The sensory system: Items | 40 |
| 3.1.2 The motor system: Actions | 40 |
| 3.1.3 The knowledge base: Schemas | 41 |
| 3.1.3.1 What is a schema? | 41 |
| 3.1.3.2 How are schemas created? The marginal attribution mechanism | 42 |
| 3.1.3.2.1 Bare schemas | 43 |
| 3.1.3.2.2 Result spinoffs | 43 |
| 3.1.3.2.3 Context spinoffs | 44 |
| 3.1.3.2.4 Conjunctive contexts and results | 45 |
| 3.1.3.2.5 Summary of marginal attribution | 46 |
| 3.1.3.2.6 Synthetic items | 46 |
| 3.2 Improving learning via focus | 48 |
| 3.2.1 The work of learning | 48 |
| 3.2.2 The basic learning algorithm | 51 |
| 3.2.3 The focused algorithm | 52 |
| 3.2.3.1 How has focus changed the computational effort? | 53 |
| 3.2.3.2 Why were these parameters chosen? | 55 |

Table of Contents

| | | |
|---|---|-----------|
| 3.3 | Results | 57 |
| 3.3.1 | Evaluation of the focused algorithm | 57 |
| 3.3.1.1 | Completeness | 57 |
| 3.3.1.1.1 | Manual evaluation methods | 58 |
| 3.3.1.1.2 | Automatic evaluation methods | 59 |
| 3.3.1.1.3 | Completeness results | 61 |
| 3.3.1.2 | Correctness | 61 |
| 3.3.2 | Comparison of the different strategies | 62 |
| Chapter 4: Goal-dependent Learning | | 67 |
| 4.1 | Introduction | 67 |
| 4.2 | Why goals? | 68 |
| 4.2.1 | Where do goals come from? | 69 |
| 4.2.2 | The relation of goals to planning | 70 |
| 4.3 | General concepts | 71 |
| 4.3.1 | Goals | 71 |
| 4.3.1.1 | The structure of goals | 72 |
| 4.3.1.2 | Using goals for learning | 76 |
| 4.3.2 | Evaluating learning | 79 |
| 4.3.2.1 | Using chaining to evaluate learning | 80 |
| 4.3.2.2 | Should we learn during evaluation? | 80 |
| 4.3.2.3 | When to plan? | 81 |
| 4.3.2.4 | How to plan? | 82 |
| 4.3.2.4.1 | Chaining | 82 |
| 4.3.2.4.2 | Path metrics | 83 |
| 4.3.2.5 | Exceptional conditions while planning | 87 |
| 4.3.2.5.1 | Actions having “no effect” | 87 |
| 4.3.2.5.2 | Serendipity | 87 |
| 4.3.2.5.3 | Goal loops | 88 |
| 4.3.2.5.4 | Other ways to be stymied | 89 |
| 4.3.2.6 | Planning for learning versus planning for evaluation | 90 |
| 4.3.2.7 | Randomizing the world | 91 |
| 4.3.2.8 | Lobotomies | 92 |
| 4.3.2.9 | Scorecards | 92 |
| 4.4 | Results | 94 |
| 4.4.1 | The work required for learning with goals | 94 |
| 4.4.2 | Changes in the characteristics of the schemas learned | 95 |
| 4.4.3 | Competence | 99 |

| | |
|--|------------|
| Chapter 5: Related Work | 111 |
| 5.1 Related Work in Machine Learning | 111 |
| 5.1.1 Introduction | 111 |
| 5.1.2 Selective attention as filtering | 113 |
| 5.1.3 Selective attention as goal-driven learning | 117 |
| 5.2 Related Work in Cognitive Science | 121 |
| 5.2.1 Introduction | 121 |
| 5.2.2 The plausibility of attention as a system of limitations | 124 |
| Chapter 6: Conclusions | 133 |
| 6.1 The effectiveness of focus of attention | 133 |
| 6.2 Future work | 134 |
| 6.2.1 Generalization and abstraction. | 134 |
| 6.2.2 Filtering | 135 |
| 6.2.3 Experimental strategy | 136 |
| 6.2.4 Goals | 137 |
| 6.2.5 Occasional defocusing | 138 |
| Appendix A: System Architecture | 139 |
| Appendix B: Implementation of Efficient Planning | 147 |
| B.1 Introduction | 147 |
| B.2 Reachability | 147 |
| B.3 Cached (lobotomized) reachable schemas. | 149 |
| B.4 Computing paths from one arbitrary schema to another | 150 |
| B.5 Computing a path between the INITIAL and FINAL sets. | 152 |
| B.6 The promise of randomized algorithms | 153 |
| Bibliography | 155 |

