

CONTENTS

PREFACE	v
ABSTRACT	vii
ACKNOWLEDGMENTS	ix
LIST OF FIGURES	xv
LIST OF TABLES	xvii
1 INTRODUCTION	1
1.1 Motivation	1
1.2 Issues	2
1.3 Contribution	5
1.4 Example of results	7
2 NATURAL OBJECT RECOGNITION	13
2.1 Visual capabilities for autonomous robots	13
2.2 Related research	14
2.2.1 Recognizing objects	15
2.2.2 Recognizing natural scenes	15
2.3 Limitations of current machine-vision technology	19
2.3.1 Shape	19
2.3.2 Universal partitioning	21
2.3.3 Contextual knowledge	23
2.3.4 Computational complexity	24
2.4 Key ideas	24
2.4.1 Context-limited vision	25
2.4.2 Global consistency	27
2.4.3 Candidate comparison to control complexity	29
2.4.4 Layered partitions	31

2.5	Experimental results	34
2.6	Conclusions	41
3	A VISION SYSTEM FOR OFF-ROAD NAVIGATION	43
3.1	Task scenario	43
3.2	Prior knowledge	44
3.3	The role of geometry	47
3.3.1	Sources and limitations of range data	47
3.3.2	Using three-dimensional geometric information	49
3.4	A vocabulary for recognition	50
3.4.1	Target vocabulary	51
3.4.2	Recognition vocabulary	51
3.5	Contextual information	54
3.5.1	Types of context	54
3.5.2	Using context	55
4	CONTEXT-BASED VISION	57
4.1	Conceptual Architecture	57
4.1.1	Overview	57
4.1.2	Context sets	60
4.1.3	Candidate generation	62
4.1.4	Clique formation	65
4.1.5	Candidate comparison	67
4.1.6	The recognition process	70
4.2	Implementation of Condor	71
4.2.1	Processing sequence	71
4.2.2	Representation of context	74
4.2.3	Context-set construction	77
4.3	Example of natural-object recognition	79
4.3.1	Candidate generation	81
4.3.2	Candidate comparison	85
4.3.3	Clique formation	89
4.3.4	Clique selection	92
4.4	Automated knowledge acquisition	95
4.4.1	Learning a description of the environment	97
4.4.2	Applying recognition procedures more effectively	98
4.4.3	Refining context sets	101
4.4.4	Discovering new procedures	105

4.5	Complexity analysis	105
4.5.1	Mathematical details	107
4.5.2	The number of cliques	108
4.6	Discussion	110
4.6.1	Reliability	110
4.6.2	Recognition strategies	111
4.6.3	Knowledge-base construction	111
4.6.4	Multiple-image interpretation	112
5	EXPERIMENTAL RESULTS	115
5.1	Evaluation scenario	115
5.2	Experimentation	119
5.2.1	Experiment 1	119
5.2.2	Experiment 2	123
5.2.3	Experiment 3	132
5.3	Analysis of results	132
5.3.1	Fixing mistakes	135
5.3.2	Accounting for success	137
5.3.3	Evaluating relevance to the goal	138
6	CONCLUSION	139
6.1	Contribution	139
6.2	Evaluation	142
6.2.1	Competence	142
6.2.2	Scalability	143
6.2.3	Generality	145
6.2.4	Evaluation plan	146
6.3	Summary	148
A	THE CORE KNOWLEDGE STRUCTURE	149
A.1	Introduction	149
A.2	Core Knowledge Structure	150
A.3	Logical Interpretation of the CKS Database	151
A.3.1	Semantics	152
A.3.2	Insertions	154
A.3.3	Queries	154
A.3.4	User-Defined Relations	158
A.3.5	Discussion	159

A.4 Slot Access	160
A.5 Summary	162
REFERENCES	165