

# Design with shape grammars and reinforcement learning

Manuela Ruiz-Montiel<sup>a</sup>, Javier Boned<sup>b</sup>, Juan Gavilanes<sup>b</sup>, Eduardo Jiménez,  
Lawrence Mandow<sup>a</sup>, José-Luis Pérez-de-la-Cruz<sup>a</sup>

<sup>a</sup>*Escuela Técnica Superior de Ingeniería Informática. Bulevar Louis Pasteur, N<sup>o</sup> 35.  
Campus de Teatinos. Universidad de Málaga. 29071 Málaga, España*

<sup>b</sup>*Escuela Técnica Superior de Arquitectura. Campus de El Ejido. Universidad de  
Málaga. 29071, Málaga, España*

---

## Abstract

Shape grammars are a powerful and appealing formalism for automatic shape generation in computer-based design systems. This paper presents a proposal complementing the generative power of shape grammars with reinforcement learning techniques. We use simple (*naive*) shape grammars capable of generating a large variety of different designs. In order to generate those designs that comply with given design requirements, the grammar is subject to a process of machine learning using reinforcement learning techniques. Based on this method, we have developed a system for architectural design, aimed at generating two-dimensional layout schemes of single-family housing units. Using relatively simple grammar rules, we learn to generate schemes that satisfy a set of requirements stated in a design guideline. Obtained results are presented and discussed.<sup>1</sup>

*Keywords:*

computational design, shape grammars, reinforcement learning, architecture

---

## 1. Introduction

The problem of computational design has been approached by Artificial Intelligence researchers since the 1960s [1, 2]. Design problem solving is

---

<sup>1</sup>This work is partially funded by: grant TIN2009-14179 (Spanish Government, Plan Nacional de I+D+i). Manuela Ruiz-Montiel is funded by the Spanish Ministry of Education through the National F.P.U. Program.

1  
2  
3  
4  
5  
6  
7  
8  
9 considered a particularly complex task, since these kind of problems are ill-  
10 structured [3], that is, they are characterized by the absence of a unique,  
11 well-defined solution. Human experts deal with these problems not only by  
12 means of knowledge about the requirements that solutions need to meet, but  
13 also with the help of creative thinking, that will lead to many valid designs  
14 substantially different between them.  
15  
16

17 In this work we deal with the computational design problem of automatic,  
18 partially-directed generation of design alternatives according to certain cri-  
19 teria. The diverse variants of this problem are of interest in many fields like  
20 architecture, engineering or packaging, where geometric design plays a cru-  
21 cial role. Certainly, an assistant system that generates and proposes different  
22 design alternatives can be quite worthy for design stakeholders in the first  
23 stages of the design process.  
24  
25

26 More specifically, we propose a computational design methodology using  
27 *shape grammars* [4, 5, 6] as design synthesis mechanism. Shape grammars  
28 have been used in the literature for numerous design tasks related to areas like  
29 architecture or industrial design. In this context, shape grammars have been  
30 used to generate designs according to functional requirements in a number  
31 of ways. However, the development of adequate methodologies to create and  
32 control the application of shape grammars remains an open research topic.  
33  
34

35 The direct, traditional approach requires a human expert to design an *ex-*  
36 *pert* shape grammar and thus hard-code the whole set of design requirements  
37 inside the rules [7, 8, 9, 10, 11, 12]. This practice is aimed at producing ad-  
38 missible solutions, but the involved grammars tend to be difficult to create,  
39 modify and maintain. Moreover, these grammars normally reduce the space  
40 of alternatives, since they have been designed with the foreknowledge of the  
41 nature of the results. Albeit this philosophy allows the quick synthesis of  
42 valid designs, expert shape grammars also facilitate obtaining designs that  
43 are prone to be very similar between them.  
44  
45

46 Alternatively, some researchers have proposed the use of more *naive* shape  
47 grammars. The blind execution of a generation system based on a naive  
48 grammar would lead to a great number of different designs, but usually the  
49 majority of them would be unfeasible. This problem has been generally  
50 addressed by means of optimization techniques that guide shape generation  
51 [13, 14, 15, 16, 17, 18, 19, 20, 21, 22]. Some of these methods yield a varied  
52 set of solutions, but we aim to go one step beyond and be able to generate  
53 every solution that meets a set of requirements.  
54  
55

56 In this paper we propose the combination of naive shape grammars with  
57  
58

1  
2  
3  
4  
5  
6  
7  
8  
9 *reinforcement learning* [23]. Instead of directly reaching solutions, the algo-  
10 rithm learns *how to* generate them; after a simple grammar is set up, our  
11 proposal incorporates a phase of *policy learning* in which a reinforcement  
12 learning algorithm is applied, yielding a policy that will guide the execution  
13 of the naive shape grammar towards feasible solutions. Design requirements  
14 are formulated as *rewards* in this phase, and the learning algorithm learns  
15 how to optimize the reward received during shape generation. This method-  
16 ology allows the generation of a large variety of feasible designs using naive  
17 rules that are easy to understand and maintain.

18  
19 We have applied the proposed methodology to a particular area of ar-  
20 chitectural design: the generation of single-family, basic house layout plans  
21 that comply with a certain guideline. This can be of great help in the first  
22 stages of architectural design, when the designer is at risk of undergoing the  
23 *blank page syndrome* and can benefit from the existence of many feasible and  
24 varied starting points that have been obtained effortlessly. These starting  
25 points are schematic, as they are not intended to be standard architectural  
26 plans, but layouts that could later work as inception for complete, standard  
27 projects carried out by an architect.

28  
29 The article is structured as follows: in Section 2 some introductory mate-  
30 rial about shape grammars and reinforcement learning is presented in order  
31 to make the paper self-contained. Then, the structure of the system is de-  
32 scribed in Section 3, addressing generation and learning issues. Next, in  
33 Section 4, results of the system are presented and analysed in terms of time  
34 generation and feasibility. In Section 5 our system is discussed from several  
35 points of view. Finally, some conclusions are drawn and possible future lines  
36 of continuation and generalization of this work are pointed out.

## 37 38 39 40 41 42 43 44 **2. Antecedents**

### 45 46 *2.1. Shape grammars*

47  
48 *Shape grammars* are a computational formalism for the generation of geo-  
49 metric shapes introduced by Stiny and Gips [4, 5, 6]. This section summarizes  
50 some basic definitions needed later in this paper. Informally, shape gram-  
51 mars are a set of rules that govern the composition of simple geometric and  
52 symbolic elements to generate complex shapes.

53  
54 A *segment* or *line*  $l = \{p_1, p_2\}$  is defined by any pair of distinct points  $p_1$   
55 and  $p_2$ , the so-called *end points* of the line. A *shape* is defined by a finite set  
56  
57  
58

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

of distinct *maximal* lines, i.e. lines that are not part of longer lines inside the shape. The representation of a shape is thus unique.

A *labelled shape*  $\sigma$  is an ordered pair  $\sigma = \langle s, P \rangle$  where  $s$  is a shape and  $P$  is a finite set of labelled points. A labelled point  $(p, A)$  is a point  $p$  with a symbol  $A$ . A labelled shape  $s_1$  is a sub-shape of another labelled shape  $s_2$  ( $s_1 \leq s_2$ ) if and only if every line and every labelled point of  $s_1$  is in  $s_2$ .

Formally, a shape grammar is a tuple  $\langle S, L, R, I \rangle$  where:

- $S$  is a finite set of shapes
- $L$  is a finite set of symbols
- $R$  is a finite set of rules  $\alpha \rightarrow \beta$ , where  $\alpha$  is a non-empty labelled shape and  $\beta$  is a labelled shape
- $I$  is a non-empty labelled shape, called *initial shape* or *axiom*.

A rule  $\alpha \rightarrow \beta$  applies to a shape  $\gamma$  when there is a transformation  $\tau$  such that  $\tau(\alpha)$  is a sub-shape of  $\gamma$ . Usually,  $\tau$  is a general geometric transformation. In this work, transformations involve translations, rotations and reflections.

The result produced by the application of a rule  $\alpha \rightarrow \beta$  to a labelled shape  $\gamma$  under transformation  $\tau$  is given by the expression  $\gamma - \tau(\alpha) + \tau(\beta)$ . This new labelled shape is obtained substituting some occurrence of  $\tau(\alpha)$  inside  $\gamma$  with  $\tau(\beta)$ .

Figure 1 displays a rule and one sample derivation, i.e., a possible sequence of shapes generated by successive applications of the rule.




Figure 1: A rule (a) and one derivation starting from an initial squared shape (b)

Shape grammars have been used in the literature for many tasks related to architectural or industrial design. In architecture, the works of Stiny

1  
2  
3  
4  
5  
6  
7  
8  
9 and Gips [8] and Cagdas [9] have managed to devise shape grammars that  
10 capture and explain the architectural style of certain buildings, and thus  
11 create designs similar to the original ones. The work of Duarte [24] provides  
12 shape grammars that generate designs according to a given housing program.  
13 In the context of industrial design, shape grammars have been used to design  
14 artefacts taking into account sets of design requirements. Some examples  
15 include the grammars proposed by Lee and Tang [25] to generate compact  
16 digital cameras, considering different combinations of form criteria; the one  
17 proposed by Agarwal, Cagan and Constantine [26] for the generation of coffee-  
18 makers, taking into account different cost-related aspects; and the one devised  
19 by MacCormack, Cagan and Vogel [12] to capture a car's brand identity.  
20 Shape grammars have also been used in the context of process planning.  
21 Brown et al. presented a method for formalising manufacturing information  
22 through shape grammars, providing semantics such that a given derivation of  
23 the grammar can be interpreted as a process plan [27]. More recently, Shea et  
24 al. developed a framework for an autonomous design-to-fabrication system  
25 that automatically fabricates customized parts [28]; in this framework, the  
26 shape rules encode primitive movements of the machine tool.

27  
28  
29  
30  
31  
32 The wide appeal of shape grammars comes on one hand from their versa-  
33 tility. It is well known that they are capable of producing any possible shape  
34 [29]. On the other hand, this formalism presents specific advantages. They  
35 provide an intuitive method for shape definition. They are also *compact*, in  
36 the sense that they can generate sophisticated and unexpected designs with  
37 just a few rules [30]. As Chase pointed out, design systems based on gram-  
38 mars have a great potential to automate and explore many design alternatives  
39 without tedious work [31]. These systems can help designers to focus on un-  
40 expected designs that otherwise could be easily overlooked. This potential is  
41 subject to the presence of a computer in charge of automatically, quickly ex-  
42 ecute shape grammars, as pencil-and-paper execution might be tedious and  
43 therefore useless for the sake of discovering many new solutions.

44  
45  
46  
47 Early shape grammars succeeded in capturing the essence of different  
48 architectural styles and helped to validate the formalism. However, the de-  
49 velopment of adequate methodologies to create and control the application  
50 of shape grammars when trying to achieve design goals remains an open re-  
51 search topic. In particular, the use of shape grammars for architectural or  
52 engineering design implies taking into consideration constraints and goals for  
53 the designed artefact.

54  
55  
56 The aforementioned issues of shape grammars can be summarized by  
57  
58

	<b>Helpful</b>	<b>Harmful</b>
<b>Internal origin</b>	<i>Strengths</i> -Versatility -Intuitive method -Compacity -Able to produce unexpected shapes -They can be automated in order to produce many design alternatives	<i>Weaknesses</i> -Creation and control of SG might be difficult when trying to achieve design goals -SG execution needs considerable computational expenses
<b>External origin</b>	<i>Opportunities</i> -General improvement in computers' capabilities -Increasing presence of computers in design processes	<i>Threats</i> -Reluctance of design practitioners to use SG for their work

Table 1: SWOT analysis of the use of shape grammars for design

means of a SWOT (*Strengths, Weaknesses, Opportunities and Threats*) analysis, used to evaluate the venture of choosing the technology of shape grammars for design tasks. A SWOT analysis involves identifying the internal and external factors that are favourable and unfavourable to achieve the objective, that is, the use of shape grammars for design (see Table 1).

This paper addresses the critical issue (in fact, one the *weaknesses* of shape grammars) of controlling the application of a set of shape grammar rules in a particular domain, so as to satisfy a set of design goals. This process is knowledge intensive, since it requires human expertise to define the appropriate set of requirements. The next sections briefly review previous research in the field. Then an innovative approach based on reinforcement learning is presented and evaluated.

## 2.2. Expert shape grammars

According to Knight, "Different approaches to connecting grammars and goals have been suggested. One approach is direct. It involves writing rules with the foreknowledge that the generated designs will meet, or start to meet,

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

given goals. In order to do this, the behaviours and outcomes of rules must be predictable in some way” [32, pag. 7].

The execution of such rules would thus easily lead to feasible designs, where feasibility is guaranteed by the structure of the shape grammar itself. We could call these *expert* shape grammars, since they include much expert design knowledge hard-coded inside their rules. There exist many references to expert grammars in the literature. For example, the work of Duarte [24] considers a set of guidelines imposed by a housing program; and also many other works aim to design architectural or engineering objects in accordance with a particular style or brand identity [8, 9, 10, 11, 12]. Figure 2 displays a fragment of the Palladian grammar used by Stiny and Mitchell [8] to generate 2-dimensional plans of villas similar to those of the famous XVI century architect Andrea Palladio (see Figure 3).




Figure 2: Seven rules of the Palladian grammar (reproduced from [8])

Some of the methods actually used to include expert knowledge inside shape grammars comprise the following:

1. Shapes in rules. This is the most direct mechanism, where desired final shapes arise from the accumulation of shapes present in the rules.
2. Control marks. Special labels added to the “real” shapes are typically used to control several aspects of rule execution, such as the order in which rules are applied, the set of rules that can be applied at a certain moment, or the way rules can be applied (i.e., transformations that can be applied to their left-hand side).




Figure 3: The Villa Malcontenta as drawn by Palladio (reproduced from [8])

A different approach is taken by Duarte in his discursive grammars [24]: each conventional shape grammar rule is accompanied by a description rule that uses extra symbolic information to give semantic to the involved shapes, thus allowing more control possibilities. Another alternative is the one taken by Agarwal, Cagan and Constantine [26], in which the intermediate shapes produced by the rules are measured by means of cost expressions. Results are delivered to an expert designer who can use them to decide which rules to apply next.

However, the knowledge engineering effort involved in the creation and modification of expert systems is important. Expert shape grammars, as well as rule-based expert systems in general, can become very difficult to create, modify and maintain. It is generally acknowledged that a “[...]mixture of knowledge types, together with the lack of adequate justifications of the different rules, makes the maintenance of such knowledge bases very difficult and time consuming” [33]. Additionally, the use of the common mechanisms to introduce expert knowledge inside rules promotes the use of deterministic shape grammars that tend to produce certain kind of shapes that are known a priori. A trade-off arises between the ability of a shape grammar to innovate and the feasibility of the produced solutions. In a correctly formulated expert shape grammar, any arbitrary execution of its rules will ideally produce feasible designs, since the shape grammar designer will have predicted in some way the nature of these outcomes. If we sacrifice the divergence capacity of



1  
2  
3  
4  
5  
6  
7  
8  
9 shape grammars in the pursuit of predictability, we are at risk of missing one  
10 of the main reasons for using this formalism as a design framework, that is,  
11 the possibility of obtaining many unforeseen, innovative solutions.  
12

### 13 14 *2.3. Naive shape grammars*

15 A different approach calls for the combination of simple shape grammars  
16 with specific methods to guide the generation process. The grammars could  
17 lead to a great diversity in results, while the control methods would guaran-  
18 tee the fulfilment of design requirements. We can call these simpler shape  
19 grammars *naive*, because an arbitrary execution of their rules without any  
20 additional guiding mechanism would not guarantee feasible designs. This  
21 approach avoids the usual difficulties in creating, modifying and maintain-  
22 ing traditional, expert shape grammars, leaving the guarantees of feasibility  
23 mainly in hands of the control mechanism.  
24

25  
26 Previous works on goal-oriented design generation with naive shape gram-  
27 mars have usually defined the task as an *optimization* problem, with a quan-  
28 tifiable objective function to maximize or minimize, and an optional set of  
29 design constraints. Two main optimization techniques have been used: *shape*  
30 *annealing* and *evolutionary algorithms*.  
31

32 *Simulated annealing* [34] is a stochastic optimization technique that em-  
33 ulates the physical process of metal annealing, i.e. intense heating and then  
34 gradual cooling until a low-energy equilibrium state is reached. *Shape an-*  
35 *nealing* [13, 15] applies simulated annealing to shape grammars. In this  
36 context, we seek a shape that minimizes the value of a given objective func-  
37 tion, interpreted as the *energy* of the shape. A set of constraints additionally  
38 determines valid transitions in the application of the shape grammar rules.  
39

40  
41 At every step of rule derivation, a rule leading to a feasible new shape is  
42 randomly selected. If the new shape obtained applying this rule has lower  
43 energy than the current one, then the transition is automatically accepted.  
44 However, if the new shape has higher energy, the transition is only accepted  
45 with a probability that depends on the energy difference and the *temperature*  
46 of the process. Initially, when the temperature is high, the process can easily  
47 escape local optima this way. As the process goes on and the temperature is  
48 reduced, the chances of escaping deeper (low-energy) optima are reduced. In  
49 an infinite process where the temperature asymptotically approaches zero,  
50 the probability of being trapped in a global minimum approaches one. In  
51 shape annealing, the stochastic sequence of rule applications goes on until the  
52 current shape cannot be further improved after a number of trials, or a limit  
53  
54  
55  
56  
57  
58

1  
2  
3  
4  
5  
6  
7  
8  
9 on the number of rule derivations is reached. The algorithm reverses rules  
10 if at a certain moment every applicable rule violates any of the constraints.  
11 The *temperature profile* that determines how quickly the process cools down  
12 is a critical parameter in simulated annealing.  
13

14 Shape annealing has been applied mainly to structural design problems  
15 such as roof truss [15, 16] or dome design [14]. The use of shape annealing  
16 can lead to a feasible near-optimal design if the algorithm parameters are  
17 correctly configured. Apart from parameter tuning, the underlying knowl-  
18 edge engineering is to a great extent very easy: we just need to compute  
19 the objective function and test whether the constraints are violated. One  
20 important drawback is that each run of the algorithm leads only to a *single*  
21 solution. If we aim to obtain many distinct designs we have to run the op-  
22 timization algorithm again and again. This approach does not learn *how to*  
23 generate designs, it just finds a solution each time it is executed.  
24

25  
26  
27 *Genetic algorithms* are an optimization technique that emulates the pro-  
28 cess of natural selection. An initial population of individuals (*phenotypes*)  
29 is represented as strings (*chromosomes* or *genotype*). Individuals are then  
30 selected and their strings combined to produce new individuals in a stochas-  
31 tic process that takes into account their fitness to a given objective function  
32 (survival of the fittest). Ideally, this iterative process produces populations  
33 with good or near-optimal individuals. Gero, Louis and Kundu [17] describe  
34 two different applications of genetic algorithms to shape grammars. In the  
35 first approach (routine design), they seek the execution order (or derivation)  
36 of the grammar rules that optimizes a given set of constraints. This can  
37 be done using rule sequences as genotypes, and evaluating the fitness of the  
38 resulting designs. In the second, more ambitious one, grammar rules them-  
39 selves are encoded for manipulation by the genetic algorithm. This way,  
40 new grammars are produced that could possibly lead to better or innovative  
41 designs. Gero and Kazakov [18] used the technique of *genetic engineering*  
42 in order to evolve shape rules by identifying rule sub-sequences that appear  
43 in the good individuals of the populations and not in the bad ones. These  
44 sub-sequences were used to make up new, complex rules that were combined  
45 with a standard genetic algorithm in order to allocate sets of tiles according  
46 to given requirements. Ang et al [19] also combined a shape grammar and  
47 a genetic algorithm in order to design Coca-Cola bottles, by means of para-  
48 metric rules and by codifying the parameters of the rule sequences inside the  
49 genotype. A similar but more complex system is described by Lee and Tang  
50 [20]. It uses genetic programming in order to determine rule parameters for  
51  
52  
53  
54  
55  
56  
57  
58

1  
2  
3  
4  
5  
6  
7  
8  
9 the problem of designing the body of a camera. Chouchoulas [21] also uses  
10 a genetic algorithm to evolve rule sequences in order to produce apartment  
11 buildings required to meet certain criteria and O’Neill et al [22] employ a  
12 genetic algorithm to design shelters.  
13

14 An advantage offered by all evolution-based approaches is simpler knowl-  
15 edge engineering. Feasibility of resulting designs is enforced by a fitness func-  
16 tion that evaluates them according to a set of criteria. Among the downsides  
17 we must mention the usual, inherent difficulty of parameter setting.  
18

19 The set of generated solutions is often restricted in the sense that it has a  
20 relatively small size compared to the diversity potential of shape grammars.  
21 Some works [20, 19] offer just a single solution to a given problem (an optimal  
22 or near-optimal one), while others [17, 18, 21] provide few distinct, feasible  
23 solutions (less than 10), which are picked from the last generation.  
24  
25

#### 26 *2.4. Reinforcement Learning*

27 In this work we aim to learn how to apply a set of given naive shape rules  
28 so as to produce feasible design alternatives. That is, we seek a function that  
29 determines a suitable *transformation* to be applied to a given shape. We will  
30 use this function at every step of the shape grammar derivation, in order to  
31 finally produce a feasible solution according to a set of design criteria. A  
32 different but related problem is that of grammatical inference [35]. However,  
33 the aim of this paper is to learn the order of rule application, and not the  
34 rules themselves.  
35  
36

37 *Reinforcement learning* [23] is an area of machine learning that deals  
38 with how to learn which actions to take in a given environment, in order  
39 to maximize a given long-term *reward*. Learning occurs through interaction  
40 with the environment, in particular, receiving positive or negative rewards  
41 after certain actions. Consider for example a baby learning to build a toy  
42 brick tower. S/he will perceive positive or negative rewards when the tower  
43 stands or falls respectively. Another example are the successful game playing  
44 programs that have acquired master levels in games like Backgammon from  
45 self-play. In this case, the only reward received from the environment is the  
46 win/draw/lose state after the game is over.  
47  
48  
49

50 Reinforcement learning has certain advantages over supervised learning  
51 (one of the most researched approaches in machine learning). Unlike super-  
52 vised techniques, reinforcement learning does not need externally provided  
53 examples; it just requires interaction with the environment. Such feature can  
54 be very useful when is impractical to obtain representative examples of all the  
55  
56  
57  
58

1  
2  
3  
4  
5  
6  
7  
8  
9 situations in which the agent is involved. Additionally, reinforcement learning  
10 allows dealing with uncertainty, making possible to consider more realistic,  
11 non-isolated environments. Due to these benefits, reinforcement learning has  
12 been successfully applied to many fields like robotics [36, 37, 38, 39, 40, 41,  
13 42, 43, 44, 45, 46], games [47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59],  
14 scheduling systems [60, 61, 62], computer vision [63, 64, 65, 66], dynamic  
15 channel allocation in cellular networks [67, 68, 69] or medicine [70, 71].  
16  
17

18 The solution to a reinforcement learning problem is a *policy*, which is  
19 basically a function that maps each possible *state* to an *action* to be taken  
20 when at that state. Independently of its immediate reward (which depends  
21 solely on the environment), each state has an associated *value*, that reflects  
22 the potential of the state for future rewards, assuming that we will follow  
23 the current policy in the future. At the same time, the rational action to  
24 take at a given state is the one that leads to a new state with maximum  
25 value. Therefore, the policy determines the value of states and vice-versa.  
26 An *optimal policy* is one that maximizes the expected value of all states and,  
27 in consequence, the expected long-term reward.  
28  
29

30 There exist three elementary methods for solving a reinforcement learning  
31 problem [23]:  
32  
33

- 34 1. *Dynamic programming*. This term gathers a collection of algorithms  
35 that can be used when a perfect *model* of the environment is given.  
36 By model we understand the transition probabilities between states  
37 (that is, the probability of reaching a state  $s'$  from a state  $s$  following  
38 the action  $a$ ) and the expected rewards of each possible transition.  
39 These methods usually *bootstrap*, that is, they update values from other  
40 previous value estimations.  
41  
42
- 43 2. *Monte Carlo* methods. Unlike the previous approach, they do not as-  
44 sume a perfect model of the environment, so they can learn through  
45 experience. They do not bootstrap, so it is necessary to wait until a  
46 final outcome (the *accumulated* reward) is available to update values.  
47  
48
- 49 3. *Temporal-Difference* learning. TD learning combines ideas from both  
50 dynamic programming and Monte Carlo methods: as the former, they  
51 learn by bootstrapping, but as the latter, they can also learn from raw  
52 experience, without a model of the environment.  
53  
54

55 Dynamic programming techniques are of limited utility because the great  
56 computational expense introduced by full sweeps over the state space. A  
57  
58

sensible option is to implement a model-free learning technique (that is, Monte Carlo or TD), that only sweeps along the states inside the experienced paths. TD( $\lambda$ ) techniques are a generalization of Monte Carlo methods and TD learning; they allow to control the level of bootstrapping by means of a parameter  $\lambda \in [0, 1]$ . If  $\lambda = 0$ , then the method is equivalent to raw TD. If higher values of  $\lambda$  are used, then we reduce the bootstrapping level, bringing the method closer to a Monte Carlo technique. *Q-learning* [72] (see Figure 4) is a TD learning algorithm with minimal convergence requirements that has been intensively researched and successfully used in many works [36, 63, 65, 38, 73, 39, 69, 54, 62, 70, 71]. Here we use the generalized version of this method, that is, Q( $\lambda$ ).

In the following we describe this technique. For simplicity, we stick to the Q(0) case, that is, the pure TD approach. In Q-learning, each pair (*state*, *action*) is associated with a long-term *value*  $Q(s, a)$ , and the policy is determined by the following rule: “in state  $s$  take the action  $a$  that yields a pair  $(s, a)$  of maximum value”, as given by the following expression,

$$\operatorname{argmax}_a Q(s, a)$$

The direct approach is to store Q-values in a table. In Q(0), When moving from state  $s$  to state  $s'$  through action  $a$ , the value  $Q(s, a)$  is updated according to the following expression,

$$Q(s, a) \leftarrow Q(s, a) + \alpha \times \delta$$

where  $\alpha \in [0, 1]$  is the so-called *learning rate* that determines speed of convergence, and  $\delta$  is the temporal difference. This is defined as follows,

$$\delta = r + \gamma \times \max_{a'} Q(s', a') - Q(s, a)$$

where  $r$  is the immediate reward at state  $s'$  (if any),  $\max_{a'} Q(s', a')$  is the maximum value currently achievable at state  $s'$ , and  $\gamma \in [0, 1]$  is the so-called *discount rate* that measures how much the agent disregards future rewards in favour of immediate ones. A value of  $\gamma = 1$  is usual when long-term future rewards are sought.

The framework of reinforcement learning adjusts well to our case: actions will be applications of shape grammar rules to the current shape. We will receive positive rewards depending on how well the resulting shape responds to certain design requirements. In our application of Q-learning to naive shape rules we have to address two practical issues.

```

1
2
3
4
5
6
7
8
9
10 Initialize  $Q(s, a)$  arbitrarily
11 Repeat (for each episode):
12   Initialize  $s$ 
13   Repeat (for each step of episode):
14     Choose  $a$  from  $s$  using policy derived from  $Q$  (e.g.,  $\epsilon$ -greedy)
15     Take action  $a$ , observe  $r, s'$ 
16      $Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a') - Q(s, a)]$ 
17      $s \leftarrow s'$ ;
18   until  $s$  is terminal
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

```

Figure 4: Q(0) algorithm (reproduced from [23])

1. The first one is related to how much we trust in our current policy during the learning stage. It is usual to follow an exploitation-exploration strategy. Most of the time, the algorithm exploits the current policy, that is, at the majority of steps it chooses the best valued action according to the table. However, from time to time, the algorithm chooses a random action in order to explore states that would otherwise never be visited by the current policy. This kind of strategy is called  $\epsilon$ -greedy, since at each step a random action is chosen with probability  $\epsilon$ .
2. The second one deals with *generalization*. In many practical cases, the number of state-action pairs is so large that cannot be stored in a table. This is the case with shape grammars, since we have as many states as possible shapes than can be generated, and as many actions as transformations can be applied to match the left-hand side of the rule to the current shape. The solution is to use a function approximator to learn the  $Q(s, a)$  values. This is done selecting a number of relevant *features*  $f_i$  to describe each pair  $(s, a)$ , and then postulating the  $Q(s, a)$  values as a function of those features.

In this work, we use a linear function of features, which worked well in our experiments over the housing unit domain. More precisely, we postulate a function,

$$Q(s, a) = \theta_1 \times f_1(s, a) + \theta_2 \times f_2(s, a) + \dots + \theta_n \times f_n(s, a)$$

where  $f_i(s, a)$  is the  $i$ -th feature of the state produced by applying action  $a$  to state  $s$ , and  $\theta_i$  is the  $i$ -th coefficient of the function  $Q$ . The coefficients  $\theta_i$  of this function are learnt at each reinforcement learning

1  
2  
3  
4  
5  
6  
7  
8  
9 step according to the following gradient-descent rule, that takes into  
10 account the temporal difference  $\delta$ ,  
11

$$12 \quad \theta_i \leftarrow \theta_i + \alpha \times \delta \times f_i(s, a)$$

13  
14  
15 The use of function approximators introduces certain complexity in Q-  
16 learning but, at the same time, the use of features to represent states  
17 has certain advantages. Many different shapes can share the same set  
18 of features, so when we update the value for a given set of features, we  
19 are learning (generalizing) over all similar pairs (*state, action*). Thus,  
20 if an adequate set of features is selected, proper Q-values can be learnt  
21 even for states that were never visited during the learning stage. The  
22 features selected for our shape grammar domain will be described in  
23 Section 3.3.  
24  
25  
26

27 The algorithm we have used for  $Q(\lambda)$  (that is, the generalized version of Q-  
28 learning, so as to control the bootstrapping level) with linear approximation  
29 and binary features can be found in [23] (p. 213).  
30  
31

### 32 **3. A system for the generation of housing unit designs**

33  
34 In this section we describe our approach combining shape grammars and  
35 reinforcement learning through a knowledge-intensive test case: the design  
36 of single-family housing units. In the conceptual stage of this design process,  
37 where the architect faces the *blank page syndrome* but also has to consider  
38 many design requirements, a system that proposes many distinct, feasible  
39 and even unexpected starting points can be of great help.  
40  
41  
42

#### 43 *3.1. Housing unit design*

44  
45 The design of housing units, either single or multi-family, is a strongly  
46 constrained process. The nature of such constraints ranges from issues such  
47 as the area of each space, to adjacency relationships between them and,  
48 of course, more detailed constraints particular to each different space. In  
49 this work we have started from the housing program proposed by the studio  
50 Montaner & Muxí [74] for the regional government of Andalusia (Spain).  
51 This program details the criteria that a basic house must fulfil depending  
52 on the number of inhabitants. By a *basic house* we understand a house  
53 that, besides satisfying some minimum habitability conditions, also offers  
54 some adaptability, that is, its spatial composition may be modified if the  
55  
56  
57  
58

1  
2  
3  
4  
5  
6  
7  
8  
9 number of inhabitants rises. The final outcome of our experiments will be  
10 two-dimensional floor distribution schemes (in the following, *schemes*) of  
11 basic, two-person housing units. The produced schemes are distributed in  
12 just one floor.  
13

14 Albeit working with floor plans does not guarantee a global approach  
15 to architectural design, it allows the establishment of relationships between  
16 spaces, the generation of circulation spaces and the definition of static places.  
17 It also allows a dimensional quantification regarding area and longitudinal  
18 measures, and therefore also concerning the house scale, making it compatible  
19 with constructive and structural dimensions that approach reality.  
20

21 In the housing proposal of Montaner, several kinds of spaces that must  
22 be present in a basic house are established. Among these we have considered  
23 three main categories: 1) specialized spaces (which need specific installa-  
24 tions), 2) non-specialized spaces (do not need specific installations, and their  
25 use is determined by the inhabitants: dining-room, living-room, bedroom)  
26 and 3) complementary spaces (such as the *distribution hall*, that allows cir-  
27 culation between spaces).  
28  
29

30 Table 2 summarizes the imposed explicit constraints. The nature of the  
31 house contour is not constrained by our guidelines [74], but in the context  
32 of a computational system we must set some additional criteria in order to  
33 avoid the generation of unfeasible contours. In our experiments we have  
34 favoured *compact* contours, (i.e. contours with a high area/perimeter ratio)  
35 over scattered ones, because the latter might make the process of establishing  
36 the internal layout of the scheme very difficult.  
37

38 In Figure 5 proximity relationships between spaces are represented graph-  
39 ically. The circle represents the external wall of the unit. Dotted lines estab-  
40 lish *possible* relations, that differ from the striped line (*adjoining* relation)  
41 in the level of importance (possible relations are less necessary than adjoin-  
42 ing ones). Our system also considers an additional constraint that is not  
43 explicitly mentioned in the housing proposal, but it is desirable to take it  
44 into account in the design of a housing unit: the entrance to the house must  
45 be near both the kitchen and the distribution hall, and not very close to the  
46 bathroom.  
47

48 Working with schemes of housing units for two people gives rise to some  
49 considerations about intimacy and privacy that have implications in the exis-  
50 tence of internal partition walls. Sharing a two-person housing unit involves a  
51 high level of intimacy between the inhabitants, and thus less privacy is needed  
52 than in the case of a higher number of people living in the house, so it is pos-  
53 sible to have a higher level of intimacy and privacy than in the case of a higher  
54 number of people living in the house, so it is possible to have a higher level of  
55 intimacy and privacy than in the case of a higher number of people living in the  
56 house, so it is possible to have a higher level of intimacy and privacy than in  
57 the case of a higher number of people living in the house, so it is possible to  
58 have a higher level of intimacy and privacy than in the case of a higher number  
59 of people living in the house, so it is possible to have a higher level of intimacy  
60 and privacy than in the case of a higher number of people living in the house,  
61 so it is possible to have a higher level of intimacy and privacy than in the case  
62 of a higher number of people living in the house, so it is possible to have a  
63 higher level of intimacy and privacy than in the case of a higher number of  
64 people living in the house, so it is possible to have a higher level of intimacy  
65 and privacy than in the case of a higher number of people living in the house,  
66



1	
2	
3	
4	
5	
6	
7	
8	
9	
10	
11	
12	
13	<b>Global requirements</b>
14	R1: Total area must be of at least $46 m^2$
15	R2: The contour must be compact (not scattered)
16	<b>Kitchen requirements</b>
17	R3: Minimum linear space must be of at least of 6 modules $60 \times 60$ cm
18	R4: Minimum distance between modules and walls: 1,10 m
19	R5: Minimum distance between modules: 1,10 m
20	<b>Bathroom requirements</b>
21	R6: At least two modules of $90 \times 180$ cm must exist
22	<b>Non-specialized spaces requirements</b>
23	R7: The area of each non-specialized space must be bigger than $9 m^2$
24	R8: A 2,8 m-diameter circle must be inscribed inside each non-specialized space
25	<b>Complementary spaces requirements</b>
26	R9: A support space that allows the circulation between spaces must exist
27	

Table 2: Requirement set for a single-family basic house (adapted from [74])

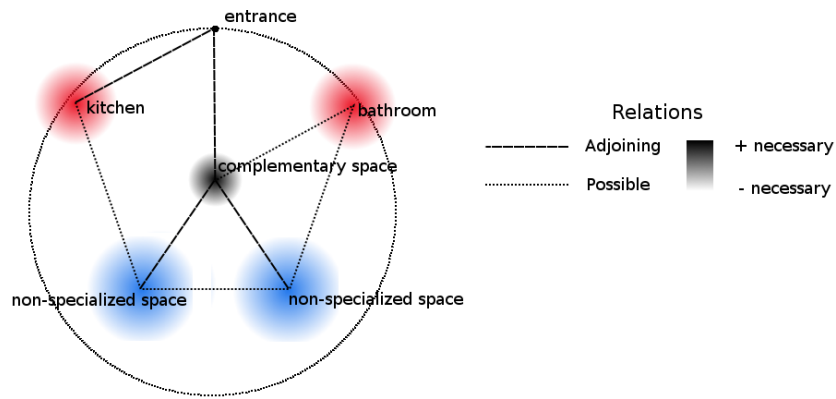


Figure 5: Proximity relationships in a single-family house (adapted from [74])

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

sible to relax the existence of internal partitions. The most necessary walls are those that isolate the bathroom from the rest of the house. Even walls isolating the kitchen are not essential. The walls that delimit non-specialized spaces such as living rooms or rooms are optional, and demarcation can be achieved by means of proper pieces of furniture.

Here we have exposed the criteria that will guide the housing unit generation process in our system. Knowledge in Table 2 and in Figure 5 will be formalized in the form of naive rules and, much more intensively, rewards, that are numerical expressions that represent the quantities to be maximized by the learning algorithm.

All the implementations along this work have been programmed using the Ruby language [75]. For visualization and geometry-related tasks we have used Trimble SketchUp [76], a 3D modelling tool that allows incorporation of Ruby scripts. The main advantages of this technological choice are that (a) SketchUp makes the edition and visualization tasks very easy, and (b) it additionally provides a complete set of geometry operations that are very useful when implementing ad-hoc shape grammars interpreters by means of Ruby. The main shortcoming of this technological choice is the increase in computing time that the use of an interpreted language may introduce.

### 3.2. Naive grammars and phases of generation

For the synthesis of housing units, we use naive shape grammars. Therefore, the creation and management of rules is to a great extent easy. However, we need additional mechanisms to control the execution of each rule, as explained in Section 2.3.

We have defined several design phases that establish an ordering in the placement of the different architectural elements. Each phase can be described as follows:

- Phase 1: Generation of a contour.
- Phase 2: Labelling the distribution hall.
- Phase 3: Placing the kitchen modules.
- Phase 4: Placing the bathroom modules.
- Phase 5: Labelling non-specialized spaces.
- Phase 6: Labelling the entrance.

1  
2  
3  
4  
5  
6  
7  
8  
9 This division in phases allows us to simplify the problem, since each phase  
10 needs to care only about the relevant set of requirements.

11 Figure 6 illustrates the grammars for phases 1-6 (there is a maximum  
12 of two rules per phase). The input to the first phase is a simple axiom: a  
13 suitably labelled module of one square meter ( $1 \text{ m}^2$ ). The shape generated  
14 in the  $n$ -th phase is the input for the  $(n + 1)$ -th phase.  
15  
16

17 In the following we explain how each rule works. Rule 1 is in charge of  
18 creating the contour. White labels are used to mark the added walls, and  
19 black labels mark the interior of the housing unit. In Figure 1 we can see an  
20 example of how this rule enlarges the small contour delimited by the axiom.  
21 Some residual walls can arise from the application of rule 1, thus we use rule  
22 2 in order to delete them. As we can see in Figure 6, the left-hand side of  
23 rule 2 ensures that it will only erase inner walls, keeping the walls delimiting  
24 the contour untouched. Rule 3 puts the label that marks the location of  
25 the distribution hall somewhere in the interior of the housing unit, which is  
26 marked with black labels. Rules 4 and 5 are in charge of placing the kitchen  
27 modules (that is, the pieces of furniture that will delimit the kitchen space).  
28 In particular, rule 4 adds the first module, which is placed right next to a  
29 contour wall. Rule 5 adds the rest of the modules, just adding one module  
30 next to an existing one. Rules 6 and 7 are similar to rules 4 and 5, but they  
31 refer to bathroom modules. Notice that the dimensions of these modules  
32 (both for the kitchen and the bathroom) are directly extracted from the  
33 housing program of Montaner (see Table 2). Rule 8 marks a non-specialized  
34 space, and it works exactly as rule 3, except for the kind of label in its right-  
35 hand side. Finally, rule 9 marks the entrance to the housing unit, placed  
36 right next to a contour wall.  
37  
38  
39  
40  
41

42 Rules are executed in order (the ordering is given by the rule numbers)  
43 and repeatedly applied as long as possible, or until a final state test has been  
44 satisfied. This final state test depends on each rule. Tests for every rule are  
45 gathered in Table 3.  
46

47 This set of rules can be used without further guidance to generate schemes;  
48 when executing each rule, all possible transformations are computed and one  
49 of them is selected at random. Two examples are depicted in Figure 11  
50 For the final presentation of the schemes, two additional operations are per-  
51 formed: (a) erasing of the labels and (b) addition of significant text in order  
52 to distinguish each space. As we will discuss in Section 4, the schemes are  
53 not actually admissible, highlighting the need of extra guiding mechanisms  
54 in order to comply with housing unit requirements.  
55  
56  
57  
58

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

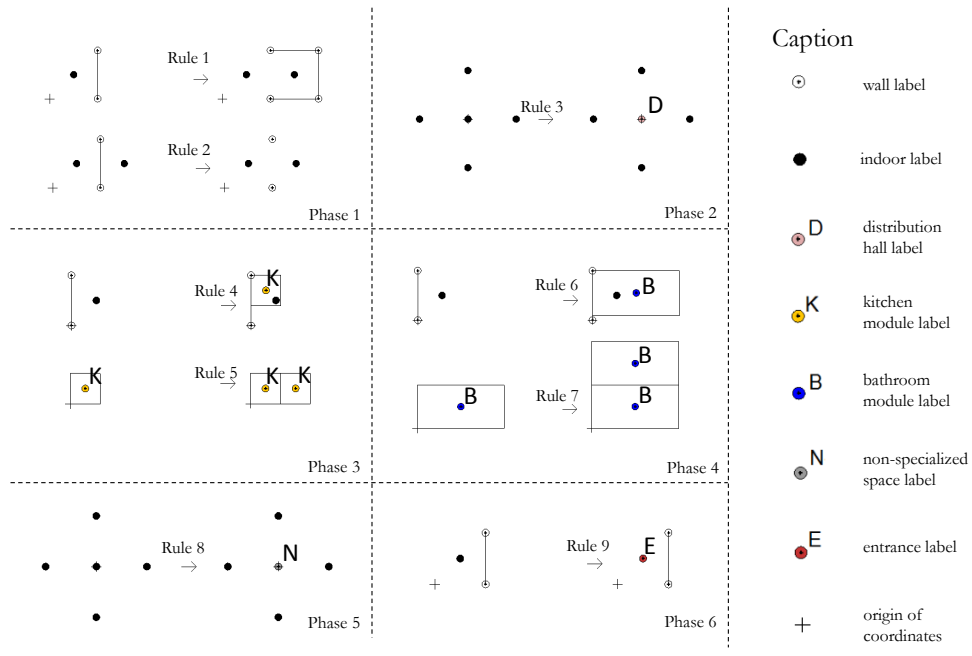


Figure 6: Naive grammars for phases 1-6

Rule	Final state test
1	Total area = 46 m <sup>2</sup>
2	Execution continues until the rule cannot be applied any longer
3	A distribution hall label exists
4	The first kitchen module is placed
5	Six kitchen modules are placed
6	The first bathroom module is placed
7	Two bathroom modules are placed
8	Two specialized labels exist
9	The entrance to the house is placed

Table 3: Final state tests for every rule




Figure 7: Effect of the application of rule 2

### 3.3. Learning processes for naive shape rules

In order to generate valid designs, we must incorporate a policy into each naive rule. These policies will select the most suitable transformation to be applied at each step in order to produce feasible solutions.

Nevertheless, not every rule in Figure 6 needs the use of a policy. For example, rule 2 cannot perform badly, since it only erases residual walls that appear from the repeated execution of rule 1 (see Figure 7).

Rule 3, which is in charge of placing the label for the distribution hall, can also perform well randomly, since when it is applied, none of the elements of the scheme has been placed yet. So this label can be placed at any point inside the scheme, and then the rest of the elements will be placed considering its situation.

Rules 1, 4, 5, 6, 7, 8 and 9 need to be guided by a policy in order to perform properly, because they have to meet requirements that have not been considered inside the rules due to their naive nature. These policies are learnt through reinforcement learning. As explained in section 2.4, in a reinforcement learning process we interact with our environment by obtaining rewards; these will be determined from the conditions described in Section 3.1, concretely the ones in Table 2 and Figure 5. We also need to establish state features, because we deal with a vast state space and thus storing the value of each state in a table is not practical; therefore, we will learn a function of the features that returns the values of states.

As an example of how policy learning is accomplished, we will consider rule number 5 to explain in detail the features and rewards that have been used. Rule 5 is in charge of placing the kitchen modules. It is depicted in figure 6. As we can see, the rule simply states that an additional module can

1  
2  
3  
4  
5  
6  
7  
8  
9 be placed next to an existing kitchen module.

10 The reward for the learning process of this rule is defined as follows: the  
11 more requirements the state complies with at a certain step, the bigger reward  
12 it gets. An advantage of dividing the process in distinct phases is that we only  
13 have to consider subsets of requirements. In this case, we have considered  
14 six requirements that are evaluated as 1 or 0 depending on whether they are  
15 fulfilled or not in the shape corresponding to a given state. The first one is a  
16 geometric constraint: every module must lie inside the contour. The others  
17 are direct translations of constraints imposed by the guideline [74]:  
18  
19  
20

- 21 •  $r_{5,1}(s) = 1$  if every module is inside the contour.
- 22
- 23 •  $r_{5,2}(s) = 1$  if every module is accessible.
- 24
- 25 •  $r_{5,3}(s) = 1$  if the distance between modules and walls is bigger than 1,1
- 26 m.
- 27
- 28 •  $r_{5,4}(s) = 1$  if the distance between non-contiguous modules is larger
- 29 than 1,1 m.
- 30
- 31 •  $r_{5,5}(s) = 1$  if the modules are at a proper distance from the distribution
- 32 hall (more than 1,2 m and less than 6).
- 33
- 34 •  $r_{5,6}(s) = 1$  if there are at least six modules.
- 35  
36  
37

38 So the reward  $r(s)$  of a state (a shape)  $s$  in the context of rule 5 is  
39 computed as the following sum:  
40

$$41 \quad r_5(s) = 3 * r_{5,1}(s) + r_{5,2}(s) + r_{5,3}(s) + r_{5,4}(s) + r_{5,5}(s) + r_{5,6}(s)$$

42  
43  
44 Some extra expert knowledge can be used to define this reward. In the  
45 case of rule 5, this expertise is translated in giving more importance to the  
46 fulfilment of the first requirement (the one that determines if the modules  
47 are inside the contour).  
48  
49

50 To define the features for this phase, we have decided to identify a feature  
51 with every single requirement, so six binary features were considered for the  
52 shapes generated by this rule. Each feature is computed over the shape  $s'$   
53 produced by applying action  $a$  to the shape of state  $s$ :  
54  
55

$$56 \quad f_{5,i}(s, a) = r_{5,i}(s'), 1 \leq i \leq 6$$

Rule	Reward
1	area/perimeter <sup>2</sup>
4	$r_{4,1}(s) + r_{4,2}(s)$
5	$3 * r_{5,1}(s) + r_{5,2}(s) + r_{5,3}(s) + r_{5,4}(s) + r_{5,5}(s) + r_{5,6}(s)$
6	$3 * r_{6,1}(s) + r_{6,2}(s) + r_{6,3}(s)$
7	$3 * r_{7,1}(s) + r_{7,2}(s) + r_{7,3}(s) + r_{7,4}(s)$
8	$3 * r_{8,1}(s) + r_{8,2}(s) + r_{8,3}(s) + r_{8,4}(s) + r_{8,5}(s) + r_{8,6}(s)$
9	$r_{9,1}(s) + r_{9,2}(s) + r_{9,3}(s) + r_{9,4}(s)$

Table 4: Rewards for every learning process

So the learned value  $Q_5(s, a)$  is determined by the function:

$$Q_5(s, a) = \theta_{5,1} \times f_{5,1}(s, a) + \dots + \theta_{5,6} \times f_{5,6}(s, a)$$

The reward  $r(s)$  refers to the short-term value of certain shape. In the end, the algorithm will learn the adequate coefficient  $\theta$  for each feature  $f$  (that is, the policy), so as to determine which features must be pursued first in order to maximize the accumulated reward obtained at the end of the rule application.

The reward expressions used for each rule are gathered in Table 4. Rewards are assigned to states at every step, except in the first rule, where they are assigned only to final states. This is due to the fact that we only have valuable information of the compactness of a contour when all the modules have been placed.

The complete set of requirements that define rewards for rules 4-9 is gathered in Table 5. All these requirements are binary, i.e one or zero when the associated predicate is satisfied or not respectively.

Table 5: Requirements for rules 4-9

---



---

<b>Rule 4</b> (placing the first kitchen module)
$r_{4,1}(s) = 1$ if the module is at a proper distance from the distribution hall (more than 2 m and less than 4)
$r_{4,2}(s) = 1$ if the module is separated from the walls by at least 1,1 m
<b>Rule 5</b> (placing the rest of the kitchen modules)
$r_{5,1}(s) = 1$ if every module is inside the contour
$r_{5,2}(s) = 1$ if every module is accessible
$r_{5,3}(s) = 1$ if the distance between modules and walls is bigger than 1,1 m

Continued on Next Page...

Table 5 – Continued

---

$r_{5,4}(s) = 1$  if the distance between non-contiguous modules is larger than 1,1 m  
 $r_{5,5}(s) = 1$  if the modules are at a proper distance from the distribution hall (more than 1,2 m and less than 6)  
 $r_{5,6}(s) = 1$  if there are at least six modules

---

**Rule 6** (placing the first bath module)  
 $r_{6,1}(s) = 1$  if the module does not overlap with kitchen modules  
 $r_{6,2}(s) = 1$  if the module is at a proper distance from the distribution hall (more than 2 m and less than 4)  
 $r_{6,3}(s) = 1$  if the module is separated enough from the walls

---

**Rule 7** (placing the rest of the bath modules)  
 $r_{7,1}(s) = 1$  if the modules do not overlap with any kitchen module  
 $r_{7,2}(s) = 1$  if the modules are at a proper distance from the distribution hall (more than 2 meters and less than 4)  
 $r_{7,3}(s) = 1$  if the modules are separated enough from the walls  
 $r_{7,4}(s) = 1$  if there are at least two modules

---

**Rule 8** (labelling non-specialized spaces)  
 $r_{8,1}(s) = 1$  if in each label a 3 meter-diameter circle can be centred, without overlapping with walls or modules  
 $r_{8,2}(s) = 1$  if there is a label separated from the bath by less than 4,5 m  
 $r_{8,3}(s) = 1$  if there is a label separated from the kitchen by less than 4,5 m  
 $r_{8,4}(s) = 1$  if there are 2 non-specialized labels separated at least by 3 m  
 $r_{8,5}(s) = 1$  if the distance from each non-specialized label to the distribution hall label is larger than 2 m and shorter than 6  
 $r_{8,6}(s) = 1$  if when there is one single non-specialized label, then there is enough space for another one

---

**Rule 9** (labelling the entrance)  
 $r_{9,1}(s) = 1$  if the entrance is separated by at least 1 m from every kitchen module  
 $r_{9,2}(s) = 1$  if entrance is less than 2 m from some kitchen module  
 $r_{9,3}(s) = 1$  if the entrance is separated by at least 4 m from every bathroom module  
 $r_{9,4}(s) = 1$  if the distance from the entrance to the distribution hall label is shorter than 4 m

---



1  
2  
3  
4  
5  
6  
7  
8  
9 Features for rules 4-9 have been defined according to the expression:  
10  $f_{i,j}(s, a) = r_{i,j}(s')$ , as we have explained for rule 5. Other approaches for  
11 features definition are possible (apart from directly translating design re-  
12 quirements). For example, in the case of rule 1, the design requirement is a  
13 high compactness factor, and we have not directly defined the features from  
14 this requirement, as in rules 4-9. Using any approach, the general intuition  
15 behind the definition of the features is that they have to be related via a  
16 linear function in order to maximize the final reward. In the case of rule  
17 1, the use of a linear function of the number of modules with 1, 2, 3 and  
18 4 neighbours makes sense, as compact shapes will probably have a certain  
19 combination of the following features:  
20  
21  
22

- 23 •  $f_{1,1}(s, a) =$  number of modules of 1  $m^2$  with 1 neighbour
- 24 •  $f_{1,2}(s, a), f_{1,3}(s, a), f_{1,4}(s, a)$  analogously to  $f_{1,1}(s, a)$  with 2, 3 and 4  
25 neighbours respectively

26  
27  
28  
29 As these features are continuous, they have been normalized between 0  
30 and 1, by dividing by the total number of modules of 1  $m^2$  that are going to  
31 exist at the end of one learning episode; in this case this number is 46.  
32

33 In the end, the learning algorithm will learn how much these features re-  
34 late to each other. If the relation were not linear, then the learning algorithm  
35 would fail. However, our results suggest linear relations capture this relation  
36 well, since scattered contours are not generated.  
37

38 Except for the first rule, for which the axiom (initial state in the learning  
39 process) is always the same (one module of 1  $m^2$ ), for the rest of the rules the  
40 initial states can be quite different. For example, in the case of rule 6, the  
41 initial state is the house contour with the label of the distribution hall and  
42 the kitchen modules already placed. The distribution hall and the kitchen  
43 modules could be in many different places, so it would be illogical to always  
44 start the learning process with the distribution hall and the modules exactly  
45 in the same place. It is better to use a set of distinct initial states for every  
46 learning process. These initial states are generated by means of the previous  
47 rules, using the already learnt policies when necessary.  
48  
49

50  
51 The coefficients for each policy were initialized to arbitrary values (in our  
52 case, they were set to zero) before the learning process. After running the  
53 reinforcement learning algorithm, the values for these coefficients are learnt,  
54 with the objective of maximizing the total accumulated reward. The learnt  
55 coefficients are gathered in table 6.  
56  
57

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

<b>Rule 1</b>
$\theta_{1,1} = 0,0037; \theta_{1,2} = 0,0044; \theta_{1,3} = 0,0338; \theta_{1,4} = 0,0624$
<b>Rule 4</b>
$\theta_{4,1} = 9,6338; \theta_{4,2} = 10,9915$
<b>Rule 5</b>
$\theta_{5,1} = 11,3647; \theta_{5,2} = 5,0703; \theta_{5,3} = 6,3998; \theta_{5,4} = 6,9184; \theta_{5,5} = 7,581; \theta_{5,6} = 1,7631$
<b>Rule 6</b>
$\theta_{6,1} = 13,9956; \theta_{6,2} = 10,9998; \theta_{6,3} = 10,7779$
<b>Rule 7</b>
$\theta_{7,1} = 6,7972; \theta_{7,2} = 7,2621; \theta_{7,3} = 6,8873; \theta_{7,4} = 6,8873$
<b>Rule 8</b>
$\theta_{8,1} = 6,3159; \theta_{8,2} = 2,9472; \theta_{8,3} = 7,2953; \theta_{8,4} = 5,1918; \theta_{8,5} = 7,4059; \theta_{8,6} = 8,3831$
<b>Rule 9</b>
$\theta_{9,1} = 7,4246; \theta_{9,2} = 2,0591; \theta_{9,3} = 5,7704; \theta_{9,4} = 1,5242$

Table 6: Policies learnt for every rule




Figure 8: Shape produced from one derivation of rule 1

For rule 1, the learnt coefficients establish a combination of the number of tiles with 1, 2, 3 or 4 neighbours. For rules 4-9, the learnt coefficients give importance to each feature regarding the maximization of the accumulated reward. Possibly not all requirements are going to be fulfilled at the end (that is, not every feature is going to be 1-valued), so the learnt coefficients give insight to determine which requirements must be pursued first in order to maximize the final total reward.

In order to clarify how the learnt policies guide rule derivation, we will go through one step of the guided application of rule 1 (the process is the same for every rule). Let us suppose that, at a certain moment, the derivation of rule 1 has yielded the shape in Figure 8.

There are many possible transformation alternatives for the next application of rule 1. Each transformation would yield a different shape, but many of these shapes share the same features. In order to represent all the shapes that share the same set of features, we use patterns according to the following method: we mark with a cross all the positions where one application of




Figure 9: Pattern that represents two shapes obtained from the shape of Figure 8 sharing the same features

rule 1 could add the extra black label of its right part (see Figure 6). For example, in Figure 9 we can see the pattern for two shapes that share the following features:

- $f_{1,1}(s, A) = 1/46$
- $f_{1,2}(s, A) = 5/46$
- $f_{1,3}(s, A) = 1/46$
- $f_{1,4}(s, A) = 2/46$

Where  $s$  is the shape depicted in Figure 8 and  $A$  is the set of transformations of rule 1 that yield some of the shapes included in the pattern of Figure 9 (in this case  $|A| = 2$ ). We can now compute the value for these features by means of the coefficients for rule 1 (see Table 6):

$$Q_1(s, A) = \theta_1 f_{1,1}(s, A) + \dots + \theta_4 f_{1,4}(s, A) = 0,1843/46$$

The application of rule 1 to the shape in Figure 8 yields four different combinations of features. In Table 7 we gather the four corresponding patterns along with their features and the computed value.

In light of the computed values gathered in Table 7, the chosen transformation is one of those that yields the pattern in the first row of Table 7 (that is, we randomly pick one action of the set  $A_1$ ), because it produces the higher value. This pattern represents the two shapes depicted in Figure 10.

When all the policies have been learnt, we can automatically produce designs with the rules guided by their policies when these are present (that is, in rules 1, 4, 5, 6, 7, 8 and 9). In Section 4 we show some obtained results.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

Pattern	Features	Value
	$f_{1,1}(s, A_1) = 0$ $f_{1,2}(s, A_1) = 4/46$ $f_{1,3}(s, A_1) = 4/46$ $f_{1,4}(s, A_1) = 1/46$	0, 2152/46
	$f_{1,1}(s, A_2) = 1/46$ $f_{1,2}(s, A_2) = 5/46$ $f_{1,3}(s, A_2) = 1/46$ $f_{1,4}(s, A_2) = 2/46$	0, 1843/46
	$f_{1,1}(s, A_3) = 1/46$ $f_{1,2}(s, A_3) = 4/46$ $f_{1,3}(s, A_3) = 3/46$ $f_{1,4}(s, A_3) = 1/46$	0, 1851/46
	$f_{1,1}(s, A_4) = 0/46$ $f_{1,2}(s, A_4) = 5/46$ $f_{1,3}(s, A_4) = 2/46$ $f_{1,4}(s, A_4) = 1/46$	0, 152/46

Table 7: Patterns yielded by one application of rule 1 to the shape in Figure 8




Figure 10: Shapes represented by the best pattern obtained from the shape of Figure 8

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

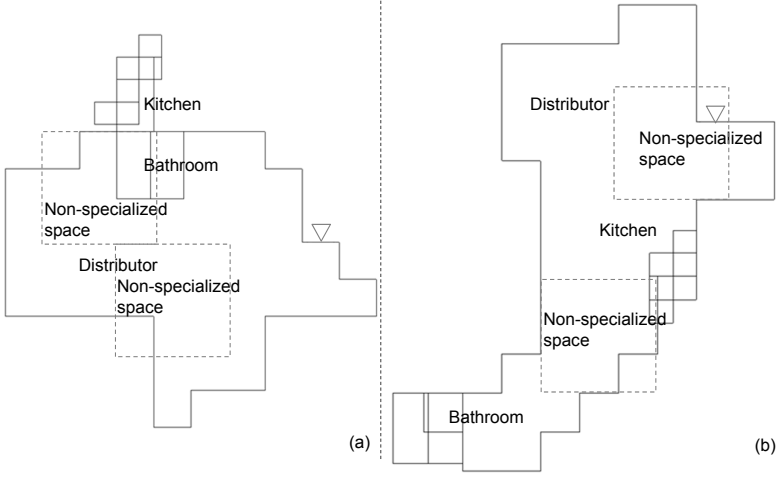


Figure 11: Two schemes obtained with naive grammars

## 4. Results

All tests reported in this work were run on an Intel Core i7 860 @2.80 GHz processor with 8GB RAM and Windows 7 (64 bits).

### 4.1. Naive grammars without policies

A number of schemes were generated with the naive set of rules depicted in Figure 6, without benefiting from any process of learning. Namely, the system was used to generate 100 schemes (two of them are shown in Figure 11). For the final presentation of the schemes we have (1) replaced labels for the kitchen, bath, distributor hall and non-specialized spaces with suitable text, (2) replaced the entrance label with an arrow, and (3) surrounded non-specialized spaces with dotted,  $3 \times 3$  m. squares.

Computation was fast. The minimum generation time for a scheme was 30,42 s and the maximum 41,88 s. The mean time was 34,87 s.

All the schemes were different, but all of them violate several requirements of the housing program described in Section 3.1. In particular, those in Figure 11 violate the following requirements of Table 2:

- R2. The contours are scattered
- R3. The kitchen modules do not form a lineal space

- R4. The distance between kitchen modules and walls is not higher than 1,1 m
- R8. A 2,8 m-diameter circle cannot be inscribed inside each non-specialized space
- R9. There is not a support space that allows the circulation between spaces.

Regarding the relevant relationships gathered in Figure 5, these schemes do not respect them mainly because a circulation cannot be established between the different spaces in the housing unit.

These results could be expected given the use of naive, non-expert shape grammars without further guidance. Grammar rules do not enforce the whole set of the involved requirements, and thus arbitrary execution is not likely to lead to feasible designs.

#### 4.2. Learning processes

In our system we have used the algorithm  $Q(\lambda)$  outlined in section 2.4 with the following parameter choice: learning rate  $\alpha = 0.1$ ; exploration rate  $\epsilon = 0.3$ ; discount rate  $\gamma = 0.8$ ; bootstrapping level  $\lambda = 0.5$ . In Section 5 we describe a parametric study that has been carried out in order to determine how different settings of parameters affect the learning processes.

In table 8 some information about the learning processes for the different rules is shown. For each rule amenable to be learnt (namely, rules 1, 4, 5, 6, 7, 8 and 9) the number of learning episodes (in Section 5 we describe how these numbers have been established) and different initial states are shown. Every learning episode ends when a final state is reached (see Table 3). The last column shows time in seconds taken by the learning process.

#### 4.3. Naive grammars with policies

A hundred designs were produced by means of generation processes guided by policies when necessary (that is, for rules 1 and 4-9 in Figure 6). Only 11 out of these 100 schemes violated some of the constraints in Table 2. Therefore, compared to the ones generated by the naive grammars alone, these are closer to fulfil the housing program described in Section 3.1.

We arbitrarily chose 12 schemes out of the valid ones. The generated schemes are depicted in Figures 12 and 13. Computation time was fast; the minimum generation time for a scheme was 33,94 s and the maximum 44,35

Rule	No. of Episodes	No. of initial states	Learning time (in seconds)
1	5000	1	4728,42
4	100	10	76,57
5	100	10	270,758
6	100	10	126,75
7	100	10	481,215
8	100	10	130,198
9	100	10	53,611

Table 8: Information of learning processes

s. The mean time was 38,46 s. The slight time increase with respect to the random execution of the naive grammars is due to the calculus of the feature values, that has to be performed in this policy-driven approach, but not in the random one. Nevertheless, the number of iterations (understood as rule derivations) is the same for both approaches, as the final step tests (shown in Table 3) are shared.

The schemes show also great design diversity. A more in-depth discussion of the quality of the results from an quantitative and architectural point of view can be found in Section 5.

## 5. Discussion

The results shown in section 4 can be discussed from several points of view. We will focus on four issues:

1. Admissibility of generated schemes
2. Architectural evaluation of the generated schemes
3. Variability of generated schemes and generation times
4. Knowledge engineering effort

*Admissibility of generated schemes.* The first goal of the system should be to generate admissible designs according to the Montaner program (that is, regarding the items in Table 2 and the relationships depicted in Figure 5). Obviously, a naive grammar generates inadmissible schemes: it is difficult for a scheme generated at random to satisfy all the requirements. It was illustrated in Section 4.1.

However, when policies are learnt the rules generate better designs, as the rewards considering Montaner criteria are maximized. Regarding conditions of Table 2 and Figure 5, we found that the requirements were violated only 22

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

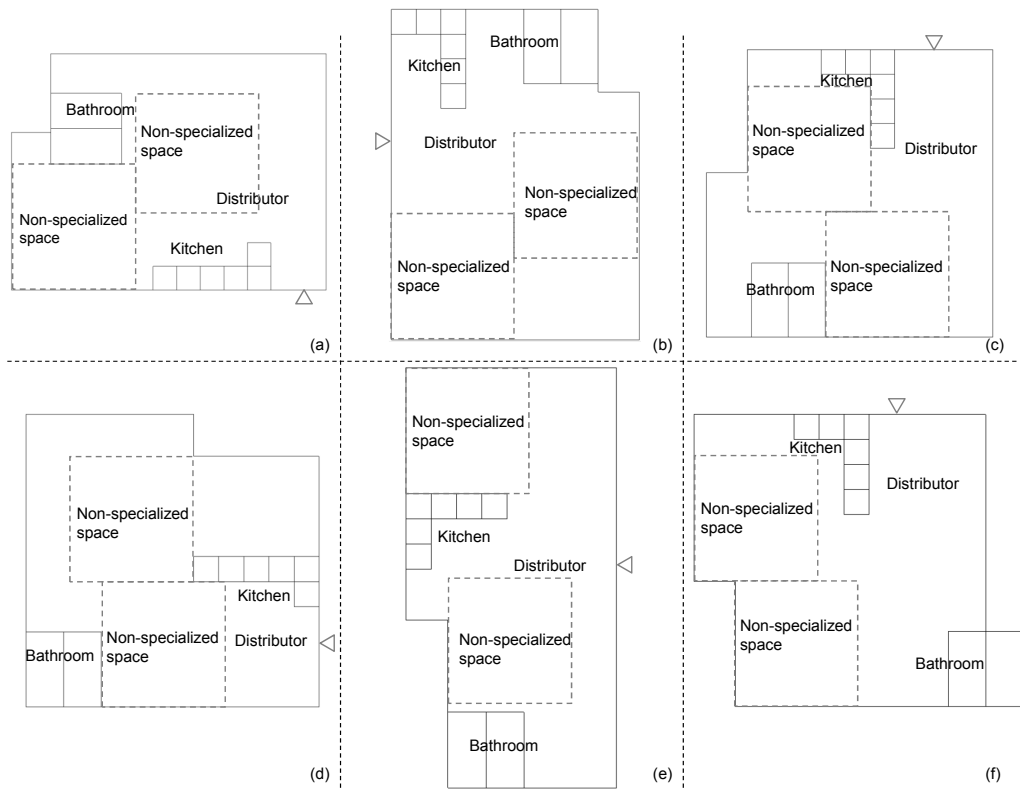


Figure 12: Some generated designs (results a-f)



1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

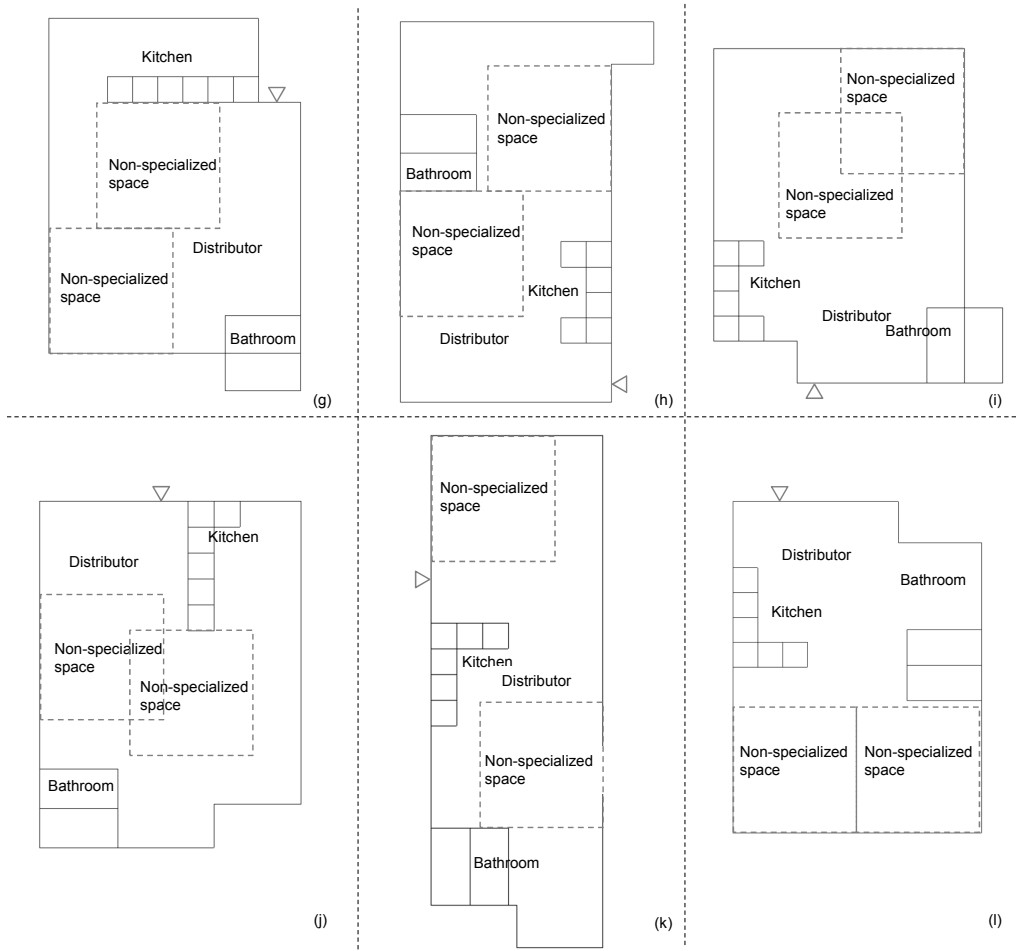


Figure 13: Some generated designs (results g-l)

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

Requirement reference	Number of schemes that violate the requirement
R3	6
R4	2
R8	5
Adjoining relation between entrance and distribution hall	3
Adjoining relation between non-specialized space and distribution hall	6
rest of requirements	0

Table 9: Requirements violated by the generated schemes

times. This fact is detailed in Table 9, where we list the number of schemes that violate each criterion.

*Architectural evaluation of the generated schemes.* Our system, applied to the generation of housing units according to a guideline, aims to help architects in overcoming the *blank page syndrome*, providing them with many distinct and feasible solutions. Once these solutions have been obtained, design practitioners can evaluate them in order to discover promising inception for their projects. This evaluation is done from a more general and subjective architectural perspective. In this section we describe this process of evaluation.

The results in Figures 12 and 13 were studied by a team of architects. From the depicted set of schemes, two were chosen by the architects as the best alternatives, and other two ones as the less adequate. Additionally, some improvements that can be easily performed in order to obtain more plausible solutions are detailed.

Schemes (a) and (b) can perform well, since non-specialized spaces are each one combined with a specialized space (kitchen or bathroom). In (a) there is an efficient combination of entrance, kitchen, distribution hall and non-specialized space that releases space and optimize the circulation. Moreover, when we enter the house we find a wide space that gives the feeling of more spaciousness. The scheme (b) aligns the kitchen and the bathroom, which is desirable given the possibility of sharing piping installations.

On the other hand, schemes (g) and (h) are not feasible from an archi-

1  
2  
3  
4  
5  
6  
7  
8  
9 tectural point of view. The kitchen and the bathroom have a diagonal dis-  
10 position that strangulates non-specialized spaces and makes the circulation  
11 difficult.  
12

13 Schemes (e) and (k) can be easily improved by changing the location of  
14 the bathroom, moving it from the bottom to the top of the scheme. Scheme  
15 (d) can be improved by moving the kitchen to fit a corner (for example, the  
16 top-right one), and thus form a more delimited space. In scheme (j) we could  
17 change the entrance position to one more centred, in order to improve the  
18 circulation inside the house.  
19

20 Interestingly, in (a) we can observe an unexpected fact concerning the  
21 entrance. Normally, this kind of housing units have the entrance more or less  
22 centred in order to make the circulation easier. Nevertheless, the entrance  
23 of scheme (a) performs very well placed in a lateral position, given the high  
24 compactness of the solution. Also, the contours of the housing units are to a  
25 great extent very unusual and inspire *grouping* or *clustering* possibilities be-  
26 tween units. These are examples of how we can find unexpected and inspiring  
27 solutions using the shape grammar framework.  
28

29 In general, except for schemes (g) and (h), all the solutions in Figures  
30 12 and 13 are feasible with minor changes or even just as they have been  
31 generated, as is the case of (a) and (b).  
32

33 It is worth questioning why the set of criteria used by the team of archi-  
34 tects has not been considered in the assessment of schemes during the  
35 learning process. Firstly, we emphasize that this architectural evaluation is  
36 part of the intended work flow inside which our system makes sense, that  
37 is: (1) generating many feasible and varied schematic designs effortlessly, (2)  
38 studying and evaluating the proposals from the point of view of the designer,  
39 and (3) using the best generated solution(s) as seed(s) for complete architec-  
40 tural projects. Secondly, the criteria considered by our team of architects do  
41 not belong to any architectural guideline (unlike the requirements that have  
42 been used in the learning process, that stem from an architectural guideline),  
43 and can be described as subjective and even elusive in some cases.  
44  
45  
46  
47  
48

49 *Variability of generated schemes and generation times.* An analysis of gen-  
50 eration time and variability of solutions is interesting in the context of an  
51 interactive system that is to be used by designers who are willing to obtain a  
52 big number of different alternatives for the first stages of the design process.  
53 These alternatives will be evaluated, and in case they do not meet design cri-  
54 teria or they just do not satisfy the designer, he or she may wish to generate  
55  
56  
57  
58

1  
2  
3  
4  
5  
6  
7  
8  
9 another big set of solutions.

10 As mentioned in section 4 and shown in Figures 12 and 13, the system  
11 generates different schemes for the given problem. In fact, when applying  
12 the policy, we choose the action that yields the maximum value. In order to  
13 produce a potentially distinct shape every time the policy is applied, we need  
14 more than a single action  $a$  to yield the maximum possible value so that we  
15 can pick one action out of the set of draws. As features aim at generalizing  
16 shapes, this behaviour is likely to occur: several actions may yield shapes  
17 with the same feature values (this fact is described in section 2.4, see Table  
18 7). Generated designs thus turn out to be substantially different in our case.

19 Once the policies are learnt, the system must run in a reasonable amount  
20 of time, suitable for an interactive environment. This goal is also achieved by  
21 the rules with the learnt policies (generation times of about half a minute per  
22 scheme, see Section 4). Notice that generation times are very similar to those  
23 for the naive grammars alone, so the effect of incorporating learnt policies to  
24 the rules is not relevant. Most time was employed in learning processes, but  
25 this is only done once for each rule.

26 The knowledge intensive domain of housing unit design will benefit from  
27 this approach, as designers will be able to consider simultaneously a huge  
28 number of automatically generated starting points (possibly unexpected) that  
29 will generally meet the considered set of design criteria.

30  
31  
32  
33  
34  
35  
36  
37 *Knowledge engineering effort.* The use of conventional expert shape gram-  
38 mars requires an important knowledge engineering effort. In the case consid-  
39 ered in Section 3.1, a set of rules that generate all valid designs of housing  
40 units according to the program under consideration would have to be de-  
41 fined. The main difficulties to overcome would be brittleness, maintenance,  
42 precedence, and unanticipated interactions of the expert shape rules.

43  
44 In contrast, the learning system proposed here makes use of naive shape  
45 grammars and an explicit rule ordering dictated by elementary considera-  
46 tions. For example, rules 4 and 5 in Figure 6 add standard kitchen modules  
47 of  $60 \times 60$  cm., but do not prescribe how these modules are to be exactly  
48 placed in the final design. These sets of naive rules are robust and easy to  
49 create, understand, modify and maintain.

50  
51 Policies on rule application are the necessary complement to this simple  
52 approach. These guarantee that rules interact in desirable ways according to  
53 the housing design program. Policies are learned automatically from rewards  
54 and features as explained in Section 3.3. Therefore, the knowledge engineer-

1  
2  
3  
4  
5  
6  
7  
8  
9 ing effort concentrates mainly on defining these rewards. In the presented  
10 case, these also follow easily from the design program.

11 However, stochastic learning algorithms, like the one used in this study,  
12 incorporate a number of parameters that need to be set before learning is  
13 actually carried out. Analysing the effect of all parameters simultaneously in  
14 algorithm performance is a combinatorially difficult task and is never carried  
15 out in practice. Instead, each parameter is usually analysed in isolation,  
16 using standard values for the others. Our case involves only four parameters,  
17 which can be summarized in order of importance as follows,  
18  
19  
20

- 21 • Bootstrapping level ( $\lambda \in [0, 1]$ ). If  $\lambda = 0$ , then the algorithm uses  
22 pure temporal differences. If  $\lambda = 1$ , then the algorithm becomes a  
23 Monte-Carlo method.  
24
- 25 • Discount rate ( $\gamma \in [0, 1]$ ). If  $\gamma = 0$ , then the algorithm favours ac-  
26 tions with immediate rewards, while with  $\gamma = 1$  long-term rewards  
27 are sought. Although our approach seeks long-term rewards, we some-  
28 what reduce this value according to the recommendations of Thrun and  
29 Schwartz [77].  
30
- 31 • Exploration rate ( $\epsilon \in [0, 1]$ ). This is the probability that the algorithm  
32 will ignore the current policy in favour of random exploration during  
33 the learning process.  
34
- 35 • Learning rate ( $\alpha \in [0, 1]$ ). This is a standard parameter in many learn-  
36 ing algorithms that controls how fast learning takes place. However,  
37 large values can lead to convergence problems.  
38  
39  
40  
41  
42

43 In the experiments described along the paper, these parameters were set  
44 to the following “standard”, conservative values:  $\lambda = 0.5, \gamma = 0.8, \epsilon = 0.3, \alpha =$   
45  $0.1$ . In order to study the influence of the parameters on the results obtained  
46 by the Q-learning algorithm, additional experiments were carried out. For the  
47 compact contour generation problem, parameters were sequentially analysed  
48 and set according to the initial values and phases described in Table 10. The  
49 first parameter analysed was the bootstrapping level ( $\lambda$ ). A wide range of five  
50 different values was analysed (0, 0.25, 0.5, 0.75, 1) while the rest of parameters  
51 were set to the values we chose ( $\gamma = 0.8, \epsilon = 0.3, \alpha = 0.1$ ). The best value  
52 for  $\lambda$  was obtained assessing the quality of the generated designs according  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

Phase	Parameter	Plausible values	Fixed values
1	$\lambda$	(0, 0.25, 0.5, 0.75, 1)	$\gamma = 0.8, \epsilon = 0.3, \alpha = 0.1$
2	$\gamma$	(0.7, 0.8, 0.9, 1)	$\lambda = \lambda_{phase_1}, \epsilon = 0.3, \alpha = 0.1$
3	$\epsilon$	(0.1, 0.2, 0.3, 0.4)	$\lambda = \lambda_{phase_1}, \gamma = \gamma_{phase_2}, \alpha = 0.1$
4	$\alpha$	(0.1, 0.15, 0.2)	$\lambda = \lambda_{phase_1}, \gamma = \gamma_{phase_2}, \epsilon = \epsilon_{phase_3}$

Table 10: Sensitivity analysis for the involved parameters in the algorithm  $Q(\lambda)$




Figure 14: Average compactness factor during the learning process

to their compactness factors. Then, the analysis of the other parameters was carried out in a similar fashion.

Most of the experimentation was devoted to parameter  $\lambda$ . Up to 3000 learning episodes were needed to produce good designs. Figure 14 shows the average compactness factor of produced designs with intermediate learned policies during the process for  $\lambda = 0.5$ . However, algorithm  $Q(\lambda)$  proved to be robust regarding parameter  $\lambda$  in this domain, returning good designs for all values. A trade-off value of  $\lambda = 0.5$  was finally chosen to take advantage of the speed provided by bootstrapping methods and the safety of Monte-Carlo techniques (see [23], Section 8.6).

Regarding the discount rate, values of  $\gamma = 0.9$  and  $\gamma = 1$  provided poor results, as already noticed by Thrun and Schwartz [77]. Therefore, the initial value of  $\gamma = 0.8$  was preserved. Finally, the algorithm proved to be robust regarding variations of  $\epsilon$  and  $\alpha$ , so the initial values were also preserved.

1  
2  
3  
4  
5  
6  
7  
8  
9 In summary, standard parameter settings proved to perform well and  
10 quite robustly in this domain.  
11

## 12 **6. Conclusions and future work**

13  
14  
15 This work deals with the computational design problem of automatic,  
16 partially-directed generation of design alternatives according to certain cri-  
17 teria. A novel proposal is presented to complement the generative power  
18 of shape grammars with reinforcement learning techniques. The approach  
19 makes use of simple (naive) shape grammars that are easy to elaborate and  
20 understand. Unlike previous approaches based on naive grammars, the gen-  
21 eration process is controlled by policies learnt with reinforcement learning  
22 techniques. The power of shape grammars allows to generate a large vari-  
23 ety of designs, while reinforcement learning helps to obtain policies that  
24 guide the generation process towards those designs that satisfy given design  
25 requirements.  
26

27  
28  
29 This methodology has been applied to generate two-person, basic housing  
30 2D scheme units. Working with floor plans simplifies geometry issues, and is  
31 enough to allow the establishment of relationships between spaces. A simple  
32 set of rules has been presented for the different phases of housing unit design.  
33 Adequate policies for design generation were learnt according to an adapted  
34 set of the constraints and conditions proposed by a design guide elaborated for  
35 the regional government of Andalusia, Spain. Learning is based on rewards  
36 and features, taken directly from the conditions established in the design  
37 guide. Learning processes are easy to define and run in reasonable time.  
38

39  
40  
41 The system generates schemes which are clearly better than the ones  
42 generated by the same grammars without learning. Schemes also present  
43 a great variability and are reasonable solutions to the housing unit design  
44 problem analysed. Time generation is also reasonable, allowing architects to  
45 consider simultaneously a huge number of automatically generated starting  
46 points (possibly unexpected) that will generally meet the considered set of  
47 design criteria.  
48

49  
50 The natural continuation of this work would be to apply the same method-  
51 ology (simple shape grammars and reinforcement learning) to more complex  
52 problems of architectural and engineering design. Our separation of the gen-  
53 eration process in sequential phases manages to reduce the set of require-  
54 ments that are taken into account at every step. However, it also introduces  
55 an important shortcoming that may have impact when dealing with more  
56  
57  
58

1  
2  
3  
4  
5  
6  
7  
8  
9 complex design problems: the learning process deals only with the local re-  
10 requirements of each phase, and thus the policies cannot reflect global issues.  
11 For example, in the context of our housing design example, a globally optimal  
12 policy would give insight about how to place the kitchen modules regarding  
13 the future placing of non-specialized spaces. A more ambitious continuation  
14 would be to relax rule ordering, allowing more flexibility in the processes of  
15 generation and learning. The application of multi-objective reinforcement  
16 learning techniques to naive design grammars is also an interesting line of  
17 future research.  
18  
19  
20  
21

## 22 **References**

- 23  
24 [1] H. A. Simon, *The sciences of the artificial*, MIT Press, Cambridge, MA.,  
25 1968.  
26  
27 [2] C. M. Eastman, *Cognitive processes and ill-defined problems: a case*  
28 *study from design*, in: *International Joint Conferences on Artificial*  
29 *Intelligence '69*, pp. 669–690.  
30  
31 [3] H. A. Simon, *The structure of ill structured problems*, *Artificial Intel-*  
32 *ligence* 4 (1973) 181–201.  
33  
34 [4] G. Stiny, J. Gips, *Shape grammars and the generative specification of*  
35 *painting and sculpture*, in: *Information Processing 71*, North-Holland,  
36 1972, pp. 1460–1465.  
37  
38 [5] G. Stiny, *Introduction to shape and shape grammars*, *Environment and*  
39 *Planning B* 7 (1980) 343–351.  
40  
41 [6] G. Stiny, *Shape. Talking about seeing and doing*, MIT Press, Cambridge,  
42 Ma., 2006.  
43  
44 [7] T. W. Knight, *Shape grammars: six types*, *Environment and Planning*  
45 *B* 26 (1999) 15–31.  
46  
47 [8] G. Stiny, W. J. Mitchell, *The Palladian grammar*, *Environment and*  
48 *Planning B* 5 (1978) 5–18.  
49  
50 [9] G. Cagdas, *A shape grammar model for designing row-houses*, *Design*  
51 *Studies* 17 (1996) 35–51.  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65



- 1  
2  
3  
4  
5  
6  
7  
8  
9 [10] U. Flemming, More than the sum of parts: the grammar of queen anne  
10 houses, *Environment and planning B* 14 (1987) 323–350.  
11  
12 [11] M. Pugliese, J. Cagan, Capturing a rebel: modeling the Harley-Davidson  
13 brand through a motorcycle shape grammar, *Research in Engineering*  
14 *Design* 13 (2002) 139–156.  
15  
16 [12] J. P. McCormack, J. Cagan, C. M. Vogel, Speaking the Buick language:  
17 capturing, understanding and exploring brand identity with shape gram-  
18 mars, *Design Studies* 25 (2004) 1–29.  
19  
20 [13] J. Cagan, W. J. Mitchell, Optimally directed shape generation by shape  
21 annealing, *Environment and Planning B: Planning and Design* 20 (1993)  
22 5–12.  
23  
24 [14] K. Shea, J. Cagan, Innovative dome design: Applying geodesic patterns  
25 with shape annealing, *Artificial intelligence for engineering design, anal-  
26 ysis and manufacturing* 11 (1997) 379–394.  
27  
28 [15] K. Shea, J. Cagan, Languages and semantics of grammatical discrete  
29 structures, *Artificial Intelligence for Engineering Design, Analysis and*  
30 *Manufacturing* 13 (1999) 241–251.  
31  
32 [16] K. Shea, J. Cagan, The design of novel roof trusses with shape annealing:  
33 assessing the ability of a computational method in aiding structural  
34 designers with varying design intent, *Design Studies* 20 (1999) 3–23.  
35  
36 [17] J. S. Gero, S. J. Louis, S. Kundu, Evolutionary learning of novel gram-  
37 mars for design improvement, *Artificial Intelligence for Engineering De-  
38 sign, Analysis and Manufacturing* 8 (1994) 83–94.  
39  
40 [18] J. S. Gero, V. A. Kazakov, Evolving building blocks for design using  
41 genetic engineering: A formal approach, in: *Advances in formal design*  
42 *mehotds for CAD*, Hall, 1996, pp. 31–50.  
43  
44 [19] M. C. Ang, H. H. Chau, A. Mckay, A. de Pennington, Combining evo-  
45 lutionary algorithms and shape grammars to generate branded product  
46 design, in: J. S. Gero (Ed.), *Design Computing and Cognition '06*,  
47 Springer Netherlands, Dordrecht, 2006, pp. 521–539.  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

- 1  
2  
3  
4  
5  
6  
7  
8  
9 [20] H. C. Lee, M. X. Tang, Evolving product form designs using parametric  
10 shape grammars integrated with genetic programming, *Artificial Intel-*  
11 *ligence for Engineering Design, Analysis and Manufacturing* 23 (2009)  
12 131–158.  
13  
14  
15 [21] O. Chouchoulas, Shape Evolution. An Algorithmic Method for Con-  
16 ceptual Architectural Design Combining Shape Grammars and Genetic  
17 Algorithms, Ph.D. thesis, University of Bath, 2003.  
18  
19 [22] M. O’Neill, J. McDermott, J. M. Swafford, J. Byrne, E. Hemberg,  
20 A. Brabazon, E. Shotton, C. McNally, M. Hemberg, Evolutionary design  
21 using grammatical evolution and shape grammars: designing a shelter,  
22 *International Journal of Design Engineering* 3 (2010).  
23  
24  
25 [23] R. S. Sutton, A. G. Barto, Reinforcement learning: an introduction,  
26 MIT Press, Cambridge, Ma., 1998.  
27  
28  
29 [24] J. P. Duarte, A discursive grammar for customizing mass housing: the  
30 case of Siza’s houses at Malagueira, *Automation in Construction* 14  
31 (2005) 265–275.  
32  
33  
34 [25] H. C. Lee, M. X. Tang, Evolving product form design using parametric  
35 shape grammars integrated with genetic programming, *Artificial Intel-*  
36 *ligence in Engineering Design, Analysis and Manufacturing* 23 (2009)  
37 131–158.  
38  
39  
40 [26] M. Agarwal, J. Cagan, K. G. Constantine, Influencing generative design  
41 through continuous evaluation: Associating costs with the coffeemaker  
42 shape grammar, *Artificial Intelligence in Engineering Design, Analysis*  
43 *and Manufacturing* 13 (1999) 253–275.  
44  
45  
46 [27] K. N. Brown, C. A. McMahon, J. H. S. Williams, Describing process  
47 plans as the formal semantics of a language of shape., *Advanced Engi-*  
48 *neering Informatics* 10 (1996) 153–169.  
49  
50  
51 [28] K. Shea, C. Ertelt, T. Gmeiner, F. Ameri, Design-to-fabrication automa-  
52 tion for the cognitive machine shop, *Advanced Engineering Informatics*  
53 24 (2010) 251–268.  
54  
55 [29] G. Stiny, Pictorial and formal aspects of shape and shape grammars and  
56 aesthetic system, Birkh auser (Basel and Stuttgart), 1975.  
57  
58

- 1  
2  
3  
4  
5  
6  
7  
8  
9 [30] P. G. Rowe, *Design Thinking*, Cambridge, Mass, MIT Press, 1987.
- 10  
11 [31] S. C. Chase, A model for user interaction in grammar-based design  
12 systems, *Automation in Construction* 11 (2002) 161–172.
- 13  
14  
15 [32] T. W. Knight, Applications in architectural design, and education and  
16 practice, Technical Report, NSF/MIT Workshop on Shape Computa-  
17 tion, 1999.
- 18  
19 [33] R. Studer, V. R. Benjamins, D. Fensel, Knowledge engineering: Princi-  
20 ples and methods, *Data and Knowledge Engineering* 25 (1998) 161–197.
- 21  
22 [34] S. Kirkpatrick, C. D. Gelatt, M. P. Vecchi, Optimization by simulated  
23 annealing, *Science*, Number 4598, 13 May 1983 220, 4598 (1983) 671–  
24 680.
- 25  
26  
27 [35] C. de la Higuera, *Grammatical Inference: Learning Automata and*  
28 *Grammars*, Cambridge University Press, New York, NY, USA, 2010.
- 29  
30 [36] P. Piggott, A. Sattar, Reinforcement learning of iterative beh-  
31aviour with multiple sensors, *Applied Intelligence* 4 (1994) 351–365.  
32 10.1007/BF00872474.
- 33  
34  
35 [37] M. J. Mataric, Reinforcement learning in the multi-robot domain, *Au-*  
36 *tonomous Robots* 4 (1997) 73–83. 10.1023/A:1008819414322.
- 37  
38 [38] C. Gaskett, L. Fletcher, A. Zelinsky, Reinforcement learning for a vision  
39 based mobile robot, in: *Intelligent Robots and Systems, 2000. (IROS*  
40 *2000)*. Proceedings. 2000 IEEE/RSJ International Conference on, vol-  
41 ume 1, pp. 403 –409 vol.1.
- 42  
43  
44 [39] W. Smart, L. Pack Kaelbling, Effective reinforcement learning for mobile  
45 robots, in: *Robotics and Automation, 2002. Proceedings. ICRA '02.*  
46 *IEEE International Conference on, volume 4, pp. 3404 – 3410 vol.4.*
- 47  
48  
49 [40] R. Tedrake, T. Zhang, H. Seung, Stochastic policy gradient reinforce-  
50 ment learning on a simple 3d biped, in: *Intelligent Robots and Systems,*  
51 *2004. (IROS 2004)*. Proceedings. 2004 IEEE/RSJ International Confer-  
52 ence on, volume 3, pp. 2849 – 2854 vol.3.
- 53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

- 1  
2  
3  
4  
5  
6  
7  
8  
9 [41] P. Abbeel, A. Coates, M. Quigley, A. Y. Ng, An application of reinforcement learning to aerobatic helicopter flight, in: In Advances in Neural Information Processing Systems 19, MIT Press, 2007, p. 2007.
- 10  
11  
12  
13 [42] G. Endo, J. Morimoto, T. Matsubara, J. Nakanishi, G. Cheng, Learning cpg-based biped locomotion with a policy gradient method: Application to a humanoid robot, *Int. J. Rob. Res.* 27 (2008) 213–228.
- 14  
15  
16  
17 [43] M. S. Erden, K. Leblebicioglu, Free gait generation with reinforcement learning for a six-legged robot, *Robotics and Autonomous Systems* 56 (2008) 199 – 212.
- 18  
19  
20  
21 [44] B. Nemeč, M. Zorko, L. Zlajpah, Learning of a ball-in-a-cup playing robot, in: *Robotics in Alpe-Adria-Danube Region (RAAD)*, 2010 IEEE 19th International Workshop on, pp. 297 –301.
- 22  
23  
24  
25 [45] J. Peters, K. Mulling, J. Kober, D. Nguyen-Tuong, O. Kromer, Towards motor skill learning for robotics, in: C. Pradalier, R. Siegwart, G. Hirzinger (Eds.), *Robotics Research*, volume 70 of *Springer Tracts in Advanced Robotics*, Springer Berlin - Heidelberg, 2011, pp. 469–482.
- 26  
27  
28  
29 [46] J. Buchli, F. Stulp, E. Theodorou, S. Schaal, Learning variable impedance control: A reinforcement learning approach, *Int. J. Rob. Res.* 30 (2011) 820–833.
- 30  
31  
32  
33 [47] G. Tesauro, Temporal difference learning and td-gammon, *Commun. ACM* 38 (1995) 58–68.
- 34  
35  
36  
37 [48] S. Thrun, Learning to play the game of chess, in: *Advances in Neural Information Processing Systems 7*, The MIT Press, 1995, pp. 1069–1076.
- 38  
39  
40  
41 [49] D. P. Bertsekas, J. N. Tsitsiklis, *Neuro-Dynamic Programming*, Athena Scientific, 1996.
- 42  
43  
44  
45 [50] J. Baxter, A. Tridgell, L. Weaver, Learning to play chess using temporal differences, *Machine Learning* 40 (2000) 243–263. 10.1023/A:1007634325138.
- 46  
47  
48  
49 [51] N. Schraudolph, P. Dayan, T. Sejnowski, Learning to evaluate go positions via temporal difference methods, in: N. Baba, L. Jain (Eds.), *Computational Intelligence in Games*, volume 62 of *Studies in Fuzziness and Soft Computing*, Springer Berlin / Heidelberg, 2001, pp. 77–98.
- 50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

- 1  
2  
3  
4  
5  
6  
7  
8  
9 [52] R. Levinson, R. Weber, Chess neighborhoods, function combination,  
10 and reinforcement learning, in: T. Marsland, I. Frank (Eds.), Computers  
11 and Games, volume 2063 of *Lecture Notes in Computer Science*, Springer  
12 Berlin / Heidelberg, 2001, pp. 133–150.  
13  
14  
15 [53] J. Schaeffer, M. Hlynka, V. Jussila, Temporal difference learning applied  
16 to a high-performance game-playing program, in: Proceedings of the  
17 17th international joint conference on Artificial intelligence - Volume 1,  
18 IJCAI'01, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA,  
19 2001, pp. 529–534.  
20  
21  
22 [54] J. Hu, M. P. Wellman, Nash q-learning for general-sum stochastic games,  
23 J. Mach. Learn. Res. 4 (2003) 1039–1069.  
24  
25  
26 [55] S. Droste, J. Fürnkranz, Learning the piece values for three chess vari-  
27 ants, International Computer Games Association Journal (2008) 209–  
28 233.  
29  
30  
31 [56] S. Gelly, D. Silver, Achieving master level play in 9x9 computer go, in:  
32 Proceedings of the 23rd national conference on Artificial intelligence -  
33 Volume 3, AAAI'08, AAAI Press, 2008, pp. 1537–1540.  
34  
35  
36 [57] C. Thiery, B. Scherrer, Building controllers for tetris (2009).  
37  
38 [58] J. Veness, D. Silver, W. T. B. Uther, A. Blair, Bootstrapping from  
39 game tree search, in: Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I.  
40 Williams, A. Culotta (Eds.), NIPS, Curran Associates, Inc., 2009, pp.  
41 1937–1945.  
42  
43 [59] I. Szita, G. Chaslot, P. Spronck, Monte-carlo tree search in settlers of  
44 catan, in: H. van den Herik, P. Spronck (Eds.), Advances in Computer  
45 Games, volume 6048 of *Lecture Notes in Computer Science*, Springer  
46 Berlin / Heidelberg, 2010, pp. 21–32.  
47  
48  
49 [60] W. Zhang, T. G. Dietterich, A reinforcement learning approach to job-  
50 shop scheduling, in: In Proceedings of the Fourteenth International  
51 Joint Conference on Artificial Intelligence, Morgan Kaufmann, 1995,  
52 pp. 1114–1120.  
53  
54  
55  
56  
57  
58

- 1  
2  
3  
4  
5  
6  
7  
8  
9 [61] M. Aydin, E. Oztemel, Dynamic job-shop scheduling using reinforcement learning agents, *Robotics and Autonomous Systems* 33 (2000) 169 – 178.  
10  
11  
12  
13  
14 [62] Y.-C. Wang, J. M. Usher, Application of reinforcement learning for agent-based production scheduling, *Engineering Applications of Artificial Intelligence* 18 (2005) 73 – 82.  
15  
16  
17  
18 [63] M. Asada, S. Noda, S. Tawaratsumida, K. Hosoda, Purposive behavior acquisition for a real robot by vision-based reinforcement learning, *Machine Learning* 23 (1996) 279–303. 10.1023/A:1018237008823.  
19  
20  
21  
22  
23 [64] J. Peng, B. Bhanu, Closed-loop object recognition using reinforcement learning, *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 20 (1998) 139 –154.  
24  
25  
26  
27  
28 [65] L. Paletta, A. Pinz, Active object recognition by view integration and reinforcement learning, *Robotics and Autonomous Systems* 31 (2000) 71 – 86.  
29  
30  
31  
32  
33 [66] J. Michels, A. Saxena, A. Y. Ng, High speed obstacle avoidance using monocular vision and reinforcement learning, in: *Proceedings of the 22nd international conference on Machine learning, ICML '05, ACM, New York, NY, USA, 2005, pp. 593–600.*  
34  
35  
36  
37  
38  
39 [67] S. Singh, D. Bertsekas, Reinforcement learning for dynamic channel allocation in cellular telephone systems, in: *In Advances in Neural Information Processing Systems: Proceedings of the 1996 Conference, MIT Press, 1997, pp. 974–980.*  
40  
41  
42  
43  
44 [68] N. Lilith, K. Dogancay, Dynamic channel allocation for mobile cellular traffic using reduced-state reinforcement learning, in: *Wireless Communications and Networking Conference, 2004. WCNC. 2004 IEEE, volume 4, pp. 2195 – 2200 Vol.4.*  
45  
46  
47  
48  
49  
50  
51 [69] S.-M. Senouci, G. Pujole, Dynamic channel assignment in cellular networks: a reinforcement learning solution, in: *Telecommunications, 2003. ICT 2003. 10th International Conference on, volume 1, pp. 302 – 309 vol.1.*  
52  
53  
54  
55  
56  
57  
58

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

- [70] Y. Zhao, M. R. Kosorok, D. Zeng, Reinforcement learning design for cancer clinical trials, *Stat Med* (2009).
- [71] R. S. H. Istepanian, N. Y. Philip, M. G. Martini, Medical qos provision based on reinforcement learning in ultrasound streaming over 3.5g wireless systems, *IEEE J.Sel. A. Commun.* 27 (2009) 566–574.
- [72] C. J. Watkins, Learning from delayed rewards, Ph.D. thesis, University of Cambridge, 1989.
- [73] J. Abounadi, D. P. Bertsekas, V. Borkar, Stochastic Approximation for Non-Expansive Maps: Application to Q-Learning Algorithms, Technical Report, *SIAM Journal on Control and Optimization*, 2000.
- [74] Montaner Muxí arquitectes, Propuesta de nueva normativa de viviendas, Technical Report, Dirección general de ordenación del territorio, Junta de Andalucía, 2008.
- [75] Ruby Language, <http://www.ruby-lang.org/>, 2011.
- [76] Trimble SketchUp, <http://sketchup.google.com/intl/es/>, 2011.
- [77] S. Thrun, A. Schwartz, Issues in using function approximation for reinforcement learning, in: *In Proceedings of the Fourth Connectionist Models Summer School*, Erlbaum, 1993.

**List of Figures**

1	A rule (a) and one derivation starting from an initial squared shape (b) . . . . .	4
2	Seven rules of the Palladian grammar (reproduced from [8]) .	7
3	The Villa Malcontenta as drawn by Palladio (reproduced from [8]) . . . . .	8
4	Q(0) algorithm (reproduced from [23]) . . . . .	14
5	Proximity relationships in a single-family house (adapted from [74]) . . . . .	17
6	Naïve grammars for phases 1-6 . . . . .	20
7	Effect of the application of rule 2 . . . . .	21
8	Shape produced from one derivation of rule 1 . . . . .	26

9	Pattern that represents two shapes obtained from the shape of Figure 8 sharing the same features . . . . .	27
10	Shapes represented by the best pattern obtained from the shape of Figure 8 . . . . .	28
11	Two schemes obtained with naive grammars . . . . .	29
12	Some generated designs (results a-f) . . . . .	32
13	Some generated designs (results g-l) . . . . .	33
14	Average compactness factor during the learning process . . . . .	38