

Reinforcement Learning for Structural Control

Bernard Adam¹ and Ian F.C. Smith, F.ASCE²

Abstract

This study focuses on improving structural control through reinforcement learning. For the purposes of this study, structural control involves controlling the shape of an active tensegrity structure. Although the learning methodology employs case-based reasoning which is often classified as supervised learning, it has evolved into reinforcement learning, since it learns from errors. Simple retrieval and adaptation functions are proposed. The retrieval function compares the response of the structure subjected to current loading event and the attributes of cases. When the response of the structure and the case attributes are similar, this case is retrieved and adapted to the current control task. The adaptation function takes into account the control quality that has been achieved by the retrieved command in order to improve subsequent commands. The algorithm provides two types of learning: reduction of control command computation time and increase of control command quality over retrieved cases. Results from experimental testing on a full-scale active tensegrity structure are presented to validate performance.

CE Database subject headings: space structures, structural control, adaptive systems, nonlinear response.

¹PhD Student, Structural Engineering Institute, Ecole Polytechnique Fédérale de Lausanne (EPFL), Station 18, 1015 Lausanne, Switzerland. E-mail: bernard.adam@a3.epfl.ch

²Professor, Structural Engineering Institute, Ecole Polytechnique Fédérale de Lausanne (EPFL), Station 18, 1015 Lausanne, Switzerland. E-mail: ian.smith@epfl.ch

Introduction

Tensegrities are spatial, reticulated and lightweight structures that are composed of struts and tendons. Stability is provided by the self-stress state between tension and compression elements. Although there is no universal definition, a recent proposal by Motro (2003) is as follows: “A tensegrity system is a system in a stable self-equilibrated state comprising a discontinuous set of compressed components within a continuum of tensioned components.” Active tensegrity structures are feasible for structures such as temporary roofs, footbridges and radio telescopes. Control of tensegrity structures has been a topic of research since the middle of the 1990s. Kawaguchi et al. (1996) studied shape and stress control of pre-stressed truss structures. However, difficulties were identified in validating numerical results through experimental testing. Djouadi et al. (1998) described a theoretical scheme to control vibrations of tensegrity systems. Sultan (1999) presented a formulation of tensegrity active control and illustrated it with the example of an aircraft motion simulator. Skelton et al. (2000) concluded that since only small amounts of energy are needed to change the shape of tensegrity structures, they are advantageous for active control. Kanchanasaratool and Williamson (2002) proposed a dynamic model to study tensegrity feedback shape control. Averseng and Crosnier (2004) studied experimentally the control of a tensegrity grid in which the actuation system is connected to the supports. Van de Wijdeven and de Jager (2005) studied an example of 2D tensegrity vibration and shape control. Few experimental studies have been designed to provide practical results. Proposals are mainly tested numerically on small, simple and symmetric tensegrity models.

The field of machine learning is divided into three sub-domains: supervised learning, unsupervised learning and reinforcement learning. The term *supervised* originates from the fact that desired outputs are provided by an external teacher (Abraham 2005). Case-based reasoning

is a method that belongs to supervised learning. Case-based reasoning systems build on findings that previous experience is useful. Humans resolve new challenges by first searching in their memory for similar tasks that they have successfully carried out in the past. Retrieved solutions are adapted to the current task (Kolodner 1993, Leak 1996). Unsupervised learning employs learning algorithms that do not require an explicit teacher. The goal of unsupervised learning is to extract an efficient internal representation of the statistical structure implicit in the inputs (Hilton and Sejnowski 1999). Reinforcement learning is based on interacting with the environment. Trial-and-error search and delayed reward are the two most important distinguishing features of reinforcement learning (Sutton and Barto, 1998; Csáji et al, 2006). Put simply, this approach involves learning from errors. Studies involving reinforcement learning for control tasks are found in the electrical and transportation engineering. For example, reinforcement learning was used for reactive power control (Vlachogiannis and Hatziaargyriou 2004). Abdulhai and Kattan (2003) describe tasks such as traffic management, infrastructure management and traffic control that were supported by reinforcement learning. In structural engineering, few examples of control procedure supported by supervised and unsupervised learning algorithms were found in the literature. Ghaboussi and Joghataie (1995) and Chen et al. (1995) proposed a supervised learning strategy to control vibrations in multi-storey buildings subjected to ground motion. The need to train this artificial neural network by an external teacher was observed to be the key drawback. Although the case-based reasoning method is widely used for tasks such as structural design, (Bailey and Smith 1994; Waheed and Adeli 2005), few studies focus on the application of case-based reasoning for structural control. An exception is Domer and Smith (2005). Madan (2005) described an unsupervised approach for training artificial neural network in active control of earthquake vibrations in building structures.

However, no studies of structural control have used the potential of reinforcement learning to improve control performance.

This paper addresses improvement of structural control through reinforcement learning. For the purposes of this study, structural control involves controlling the shape of an active tensegrity structure. The proposed learning algorithm decreases control command computation time and improves control command quality over retrieved cases. When a case is retrieved, adaptation is direct. The adaptation function is based on a local linear-elastic assumption. Moreover, quality of retrieved control commands is involved in adaptation to improve control quality over time. Previous studies at EPFL are described in the next section. Notations, quantities and learning methodology are reviewed in the following section. The results section provides an experimental validation of the proposed methodology. The paper ends with a discussion of future work.

Previous studies at EPFL

Research into active structures has been carried out at EPFL since 1996. Shea and Smith (1998) proposed that the ultimate goal of *intelligent structures* is to maintain and improve structural performance by recognizing changes in behaviors and loads, adapting the structure to meet performance goals, and using past events to improve future performance. Shea et al. (2002) presented a computational framework based on intelligent control methodology that combines reasoning from explicit knowledge, search, learning and planning to illustrate a vision of intelligently controlled civil structures. Fest (2002) provided a description of an active tensegrity structure and control system (Figure 1). It covers a surface area of 15 m^2 , has a static height of 1.20 m and withstands a distributed dead load of 300 N/m^2 . It is composed of 30 struts and 120

tendons. Struts are fiber reinforced polymer tubes of 60mm diameter and 703 mm² cross section. Tendons are stainless steel cables of 6 mm in diameter. The structure is equipped with an active control system: ten active struts and three displacement sensors. The ten active struts are placed in in-line pair in each module. They allow for strut length adjustment. Displacement sensors measure vertical displacements at three top surface nodes. Figure 2 provides a view of the structure from the top. Support conditions are specified with arrows. Top surface nodes are numbered. Active struts are numbered from 1 to 10. The three displacement sensors are placed at three nodes of the top surface: 37, 43 and 48. Shape control has involved satisfying a serviceability objective: maintaining the slope of the top surface of the structure when the structure was subjected to a loading. Top surface slope S is calculated from the vertical coordinates of the top surface three nodes, z_{37} , z_{43} and z_{48} , that are equipped with displacement sensors (Figure 2). It is determined as follows:

$$S = \frac{1}{L} \left(z_{43} - \frac{z_{37} + z_{48}}{2} \right) \quad (1)$$

where L is the horizontal distance between nodes 43 and middle of segment 37 – 48. The most challenging part of the study was the computation of control commands that modify the self-stress state in order to satisfy this objective (Fest et al. 2003). Since the structure was equipped with three displacement sensors and ten active struts for control flexibility, there was no closed-form solution for active strut adjustments given a required slope. A stochastic search algorithm (Raphael and Smith, 2003) was selected as the best method to accommodate the exponentially complex generate-test process that was needed to compute control commands (Domer et al. 2003, Fest et al. 2004). However, control command computation was time consuming. Moreover, in this study, load values and their location were known input values.

Domer and Smith (2005) studied the capacity of the structure and its control system to reduce control command computation time. Case-based reasoning was used to speed up the process. When the structure was subjected to loading, a similar configuration was retrieved from the case base and its control command was adapted to the current task. As more cases were added to the case-base, the average time necessary to compute control commands decreased. Domer (2003) showed that search time can decrease from approximately one hour down to one to two minutes. Thus, the structure was observed to learn from its experience. Clustering cases in the case-base and maintenance strategies were proposed to achieve retrieval efficiency. In addition, an artificial neural network was used to correct for inaccuracies in the structural model such as joint friction (Domer and Smith 2005). Through correcting the numerical model with a neural network, accuracy of predictions was enhanced. This study concluded that structural performance could be improved by judicious combinations of advanced algorithms. Although control command computation time decreased, control command quality enhancement was not studied.

Adam and Smith (2007) proposed a multi-objective approach to support tensegrity shape control. Since more robust control commands were computed using this approach than a single objective method, the structure was observed to accommodate multiple loading events over its service life more robustly than with single objective strategies. However, this methodology was still time consuming (computation times of 45min using a personal computer: Pentium 4, 3.08 GHz, 1.00GB of RAM). Nonlinear effects of control command application were also observed to alter control quality in some situations.

Adam and Smith (2006) described methodologies for self diagnosis in order to control an active tensegrity structure in situations of partially defined loading events and damage. Self-diagnosis solutions were successfully used to compute control commands for shape control and self repair.

Self-repair abilities were demonstrated experimentally. Self-repairing control commands could increase stiffness and decrease stresses by modifying the self-stress state of a damaged structure. Nevertheless, control command computation was time consuming. Moreover, nonlinear effects of control command applications were not overcome and this was observed to alter control quality in some situations.

Glossary

Notations and quantities that are specific to this study are defined in this section. They will be used throughout this paper. When subjected to loading, the structure displaces. The vertical coordinate of node i in deformed configuration is indicated by z_i' . Consequently the top surface slope after loading S' is calculated as follows (Figure 2):

$$S' = \frac{1}{L} \left(z_{43}' - \frac{z_{37}' + z_{48}'}{2} \right) \quad (2)$$

Slope unit is mm/100m. The top surface slope deviation $TSSD$ is determined using equations (1) and (2) as follows:

$$TSSD = S' - S \quad (3)$$

Evaluation of the response of the structure to a loading event involves active control perturbations of the structure through 1mm elongation of active struts. For a 1mm elongation of a particular active strut, vertical coordinate of node i is labeled z_i'' and top surface slope after perturbation S'' is expressed as follows:

$$S'' = \frac{1}{L} \left(z_{43}'' - \frac{z_{37}'' + z_{48}''}{2} \right) \quad (4)$$

The resulting slope variation ΔS has the following expression using equations (2) and (4). It describes the influence of a particular active strut on the top surface slope:

$$\Delta S = S'' - S' \quad (5)$$

$\Delta S(j)$ is a measure of slope variation resulting from a 1mm elongation of active strut j . The influence vector \mathbf{v} contains the 10 slope variations ΔS that correspond to the 10 active struts.

According to equation (5), its expression is:

$$\mathbf{v} = [\Delta S(1) \quad \Delta S(2) \quad \dots \quad \Delta S(10)]^T \quad (6)$$

When top surface slope has been corrected through a control command application, the vertical coordinate of node i is labeled z_i''' and corrected top surface slope S''' is calculated as follows:

$$S''' = \frac{1}{L} \left(z_{43}''' - \frac{z_{37}''' + z_{48}'''}{2} \right) \quad (7)$$

The slope compensation SC is defined in order to evaluate top surface slope correction quality using equations (1), (2) and (7). Its expression is:

$$SC = \frac{S' - S'''}{S' - S} \quad (8)$$

In order to evaluate the transversal behavior of the structure, transversal slope deviation quantity TSD is introduced (Figure 2). Its expression is:

$$TSD = \frac{1}{L_t} [(z'_{48} - z'_{37}) - (z_{48} - z_{37})] \quad (9)$$

where L_t is the horizontal distance between nodes 37 and 48 (Figure 2).

Learning methodology

Although the learning methodology employs case-based reasoning, which is often classified as supervised learning, it has evolved into reinforcement learning since it learns from errors in order to support incremental improvements of control commands. A case is composed of the following components: case attributes, control command, slope compensation (Table 1). Case attributes are the response of the structure in terms of top surface slope deviation, influence vector and transversal slope deviation. The control command is composed of ten active strut elongations and contractions. The last component is the slope compensation that is measured on the structure when the control command has been applied.

Intelligent control methodologies such as self diagnosis, multi-objective control and reinforcement learning are presented within the control loop in Figure 3. Retrieval, adaptation and memorization are direct. When a case is retrieved, no iteration is needed for adaptation. These operations improve control quality and speed up control command computation using retrieved cases.

Retrieval

It was observed in Adam and Smith (2006), that loads of similar magnitudes and in nearby locations induce similar responses. The retrieval function is based on this property. The response of the structure to the current load is evaluated in terms of top surface slope deviation (3), influence vector (6) and transversal slope deviation (9). This response is compared with the attributes of the cases. Similarity conditions are defined according to experimental testing results of load identification (Adam and Smith, 2006). The thresholds correspond to the accuracy of the

active control system. Within these limits the behavior has been observed to be similar. Similarity conditions are described below:

$$|TSSD_{case} - TSSD_m| \leq 2.2 \text{ mm}/100m \quad (10)$$

$$\frac{TSD_{case}}{|TSD_{case}|} = \frac{TSD_m}{|TSD_m|} \quad (11)$$

$$|\mathbf{v}_{case} - \mathbf{v}_m| = \sqrt{\sum_{j=1}^{10} (\Delta S_{case}(j) - \Delta S_m(j))^2} \leq 1.1 \text{ mm}/100m \quad (12)$$

where the subscript *case* refers to case attributes and the subscript *m* to measured values. When conditions (10), (11) and (12) are true for a particular case, the behavior of the structure subjected to the current load is said to be similar to this case. This case is retrieved. In situations where similarity conditions are true for more than one single case, the closest case according to (12) is retrieved.

Adaptation

Once a case is retrieved, the control command is adapted for shape control of the structure subjected to current loading. For the purpose of this study, a simple adaptation function based on a locally elastic-linear assumption is proposed. While this structure behaves in a geometrically nonlinear manner, this assumption is tested for a problem specific case using measured values.

The adaptation function is written as follows:

$$\mathbf{x} = \frac{TSSD_m}{TSSD_{case} \cdot SC_m(\mathbf{x}_{case})} \mathbf{x}_{case} \quad (13)$$

Control commands are scaled with respect to the ratio $TSSD_m/TSSD_{case}$ and the experimentally observed slope compensation $SC_m(\mathbf{x}_{case})$. When control command \mathbf{x} is applied, the

experimentally observed slope compensation $SC_m(\mathbf{x})$ is measured. It is rarely equal to 100%. Since the proposed learning algorithm takes into account SC_m , this algorithm learns from errors.

Memorization

In situations where the control command that is adapted from a retrieved case improves slope compensation, the current case replaces the retrieved case. This is reinforcement learning. In situations where the control command is computed using multi-objective search, a new case is memorized.

Experimental results

The proposed methodology has been tested experimentally; the learning process was observed over 80 randomly generated loading and subsequent control events. Single point loads ranging from 190N to 1200N were applied to the top surface nodes.

Control improvement over three retrieved cases

Figure 4 shows how slope compensation improves over 3 of the 80 loading events. Although these three loading events A (379N at node 51), B (410N at node 51) and C (444N at node 51) are different, they brought about a similar response of the structure. The horizontal axis describes steps of 1mm active strut length adjustments and the vertical axis is the normalized top surface slope deviation. The zero value on vertical axis is the initial top surface slope.

Load A is applied to the structure. Since the response of the structure is not similar to attributes of any case, load identification is carried out (Adam and Smith, 2006). Control command A is computed through multi-objective search (Adam and Smith, 2007) according to the load

identification solution. Multi-objective search requires 1500 iterations of ParetoPGSL (45min using a personal computer: Pentium 4, 3.08 GHz, 1.00GB of RAM). When control command A is applied to the structure, the slope compensation is observed to reach 86%. This control event creates a new case which is called case A.

Load B is applied to the structure. Since the loading event B induces a response of the structure that is similar to the attributes of case A, control command B is adapted from control command A. Adaptation is direct and no iteration is needed. When control command B is applied to the structure, slope compensation is observed to improve with respect to slope compensation of control event A. It is equal to 103%. Case B replaces case A.

Load C is applied to the structure. Loading event C induces a response of the structure that is similar to the attributes of case B. Therefore, control command B is adapted in order to control the structure subjected to loading event C. Control command adaptation is direct. When control command C is applied to the structure, slope compensation is once again observed to improve; it is equal to 98%. Case C replaces case B.

The following observations are of interest:

- Slope compensation of the current control event generally improves over retrieved cases.
- While control command computation takes about 45min with multi-objective search, retrieval and adaptation takes around 0.5s using a personal computer: Pentium 4, 3.08 GHz, 1.00GB of RAM.
- Since new cases replace retrieved cases, case maintenance is assured in the case base.
- Through these retrieved cases, this methodology can be seen as an iterative method that supports a nonlinear task.

- While the risk of cycling cases exists, it is not expected to alter learning efficiency since cases are replaced by cases that are observed to be of better quality.

Reinforcement learning over 80 loading and control events

Figure 5 plots the average error between the experimentally observed slope compensation and 100% (vertical axis) over 80 loading and control events (horizontal axis). The average error is calculated over ten control events. A general decrease of the error is observed over 80 loading and control events. Learning is not a monotonic process. Control event set 61 to 70 shows a notable decrease in control quality. This observation is explained by various effects:

- In situations where no case is retrieved due to the random nature of learning processes, performance momentarily decreases.
- Since tensegrity structures behave in a geometrically nonlinear manner, the response of the structure to the same load at time t and t' can be different when several loading and control events have occurred between time t and t' . Loading and control history induce deformations that influence the response of geometrically nonlinear structures. Therefore, cases are not retrieved and performance momentarily decreases.
- In situations where the adapted control command overstresses the structure, self diagnosis is performed and a multi-objective control command is computed.
- When cases are replaced, case attributes are also replaced. This might lead to not retrieving cases in situations where a case would have been retrieved prior to replacement.

Figure 6 plots the average number of iterations for ten control command computations (vertical axis) versus 80 loading and control events (horizontal axis). A simultaneous effect is that the

average number of iterations to compute ten control commands generally decreases over 80 loading and control events. Once again the observation is made that over 80 random loading events, learning is not monotonic. Figure 6 shows that average control command computing time increases for loading event sets 31-40, 51-60 and 71-80. The reasons are, again, related to the random nature of learning processes and the geometrical nonlinear behavior of tensegrity structures, overstressing control commands and loss of information during case replacement.

Number of cases in the case base

We examine now another aspect of this learning process. The effectiveness of the learning algorithm depends on the evolution of the number of cases in the case base. Figure 7 shows in black dots this evolution over the 80 loading and control events. Since new cases of better quality replace retrieved cases, the derivative of this evolution decreases over 80 control events. This evolution can be approximated by a second order polynomial function:

$$Nb_{cases} = -\frac{1}{320}x^2 + \frac{37}{40}x \quad (14)$$

An extrapolation of function (14) beyond observed values is drawn in gray in Figure 7. According to this expression, case-base saturation is obtained after about 150 loading and control events. The case base is saturated when about 70 cases are memorized. In other words, after saturation, cases are retrieved at each loading event. However, more loading and control events would be required to fully validate this affirmation.

Observations

The learning process is observed over 80 random loading and control events. The following observations are of interest:

- The methodologies for retrieval and adaptation generally allow for control quality and control rapidity improvement over retrieved cases.
- Since current cases replace retrieved cases when control quality improves, better cases are gradually stored and naturally classified in the case base. However, case attributes of replaced cases are also replaced. This leads to a loss of information.
- The reinforcement learning algorithm can be seen as an iterative method over retrieved cases that successfully accommodates nonlinear behavior.
- Geometrically nonlinear behavior of tensegrity structures can momentarily slow down the learning process.
- Control performance may decrease momentarily. Since control is performed over random loading events, cases may momentarily not be retrieved.
- The derivative of the evolution of the number of cases in the case base decreases as retrieved cases are replaced by new cases over 80 loading and control events.
- It is expected that after about 150 random loading events, cases are retrieved for each control event.

Future work

Extension of the reinforcement learning algorithm is needed to be usable in practical situations. This algorithm should be useful for scenarios of multiple loading events and damage. Further study should look into the growth of the case base to identify a saturation point for more complex loading and damage as well as for scenarios of multiple events. Studying the efficiency of this learning methodology for more complex structures would also be of interest.

Conclusions

Since the learning algorithm learns from errors, errors decrease as cases are retrieved and adapted. The following conclusions are drawn:

- The learning algorithm allows for two types of learning: decreased time for control command computation and increased control command quality as cases are retrieved and adapted.
- Success of these learning types is related to the formulation of case retrieval and adaptation, as well as the number of cases in the case base.
- Since cases are classified and iteratively replaced in the case base, case-base management, such as clustering, is not needed.
- Since the case base is expected to reach a saturation point where a case is retrieved for each control event, the need for a very large number of cases for learning algorithm efficiency is not expected.
- Although case-based reasoning is normally classified as supervised learning, the proposed learning methodology proposed in this work contributes to reinforcement learning.
- Interaction between learning algorithms and sensor devices is attractive for improving control tasks.

More generally, it is also demonstrated that reinforcement learning has the potential to improve active control of tensegrity structures and this could evolve toward more types of adaptive civil structures.

Acknowledgments

The authors would like to thank the Swiss National Foundation for supporting this work. B. Domer and P. Kripakaran are thanked for discussions and advice. E. Fest built the structure and the control system. B. Domer subsequently improved control. B. Raphael provided support during programming of the control system. We are also grateful to Passera & Pedretti SA (Lugano, Switzerland), Lust-Tec GmbH (Zürich, Switzerland) and P. Gallay for their contributions.

Notations

S	Top surface slope of the structure at initial state
S'	Top surface slope after loading
S''	Top surface slope after perturbation
S'''	Top surface slope after correction
SC	Slope compensation
TSD	Transversal slope deviation
$TSSD$	The top surface slope deviation
ΔS	Slope variation
v	Influence vector
x	Control command

References

- Abdulhai, B. and Kattan, L., (2003), "Reinforcement learning: Introduction to theory and potential for transport applications¹", *Can. J. Civ. Eng.*, 30, 981-991.
- Abraham, A., (2005), "Nature and Scope of AI Techniques", *Handbook of Measuring System Design*, edited by Peter H. Sydenham and Richard Thorn, John Wiley & Sons, Ltd.
- Adam, B. and Smith, I.F.C., (2007), "Tensegrity Active Control: a Multi-Objective Approach", *Journal of Computing in Civil Engineering*, Vol 21, No 1, 2007, pp 3-10.
- Adam, B. and Smith, I.F.C., (2007) "Self Diagnosis and Self Repair of an Active Tensegrity Structure", *Journal of Structural Engineering*, accepted.
- Averseng, J., and Crosnier, B. (2004). "Static and Dynamic Robust Control of Tensegrity Systems", *Journal of The International Association for Shell and Spatial Structures*, 45, 169-174.
- Bailey, S. and Smith, I.F.C., (1994), "Case-based preliminary building design", *Journal of Computing in Civil Engineering*, ASCE, 8, No. 4, 454-68.
- Chen, HM, Tsai, KH, Qi, GZ, Yang, JCS, and Amini, F., (1995), "Neural network for structure control", *Journal of Computing in Civil Engineering*, 9(2), 168-176.
- Csáji, B.C. Monostori, L. and Kádár, B. (2006) "Reinforcement learning in a distributed market-based production control system" *Advanced Engineering Informatics*, 20(3), pp279-288
- Djouadi, S., Motro, R., Pons, J.C., and Crosnier, B., (1998), "Active Control of Tensegrity Systems", *Journal of Aerospace Engineering*, 11, 37-44.
- Domer, B., and Smith, I.F.C., (2005) "An Active Structure that learns", *Journal of Computing in Civil Engineering*, 19(1), 16-24.

Domer, B., Raphael, B., Shea, K. and Smith, I.F.C. (2003), “A study of two stochastic search methods for structural control”, *Journal of Computing in Civil Engineering*, 17(3), 132-141.

Domer, B., (2003). “Performance enhancement of active structures during service lives”, Thèse no 2750, Ecole Polytechnique Fédérale de Lausanne.

Fest, E., Shea, K., and Smith, I.F.C., (2004), “Active Tensegrity Structure”, *Journal of Structural Engineering*, 130(10), 1454-1465.

Fest, E., Shea, K., Domer, B, and Smith, I.F.C., (2003), “Adjustable Tensegrity Structure”, *Journal of Structural Engineering*, 129(4), 515-526.

Fest, E., (2002). “Une structure active de type tensegrité”, Thèse no 2701, Ecole Polytechnique Fédérale de Lausanne.

Ghaboussi, J. and Joghataie, A., (1995), “Active Control of Structures Using Neural Networks”, *Journal of Engineering Mechanics*, Vol. 121, No. 4, 555-567.

Hilton, G. and Sejnowski, T.J., (1999), “Unsupervised Learning”, MIT Press.

Kanchanasaratool, N., and Williamson, D., (2002), “Modelling and control of class NSP tensegrity structures”, *International Journal of Control*, 75(2), 123-139.

Kawaguchi, K., Pellegrino, S. and Furuya, H., (1996), “Shape and Stress Control Analysis of Prestressed Truss Structures”, *Journal of Reinforced Plastics and Composites*, 15, 1226-1236.

Kolodner, J.L. (1993). “Case-Based Reasoning”, Morgan Kaufmann Publishers Inc., San Mateo, CA.

Leake, D.B. (1996), *Case-based reasoning: Experiences, lessons, & future directions*, D.B. Leake, ed., California Press, Menlo Park., Calif.

- Madan, A., (2005), "Vibration control of building structures using self-organizing and self-learning neural networks", *Journal of sound and Vibration*, 287, 759-784.
- Motro, R. (2003). "Tensegrity: Structural systems for the future", Hermes, ed., Penton Science, U.K., 280.
- Raphael, B. and Smith, I.F.C., (2003) "A direct stochastic algorithm for global search", *J of Applied Mathematics and Computation*, 146(2-3), 729-758.
- Shea, K., and Smith, I.F.C. (1998). "Intelligent structures: A new direction in structural control", *Artificial Intelligence in Structural Engineering, Computer Science, LNAI 1454*, Springer, Heidelberg, pp 398-410.
- Shea, K., Fest, E. and Smith, I.F.C. (2002) "Developing intelligent tensegrity structures with stochastic search" *Advanced Engineering Informatics*, Vol 16, No 1, pp 21-40.
- Skelton, R.E., Helton, J.W., Adhikari, R., Pinaud, J.P. and Chan, W., (2000), "An introduction to the mechanics of tensegrity structures", *Handbook on mechanical systems design*, CRC, Boca Raton, Fla.
- Sultan, C., (1999) "Modeling, design and control of tensegrity structures with applications", PhD thesis, Purdue Univ., West Lafayette, Ind.
- Sutton, R. and Barto, A., (1998), "Reinforcement learning: an introduction", MIT Press, Cambridge Mass.
- Vlachogiannis, J.G. and Hatziargyriou, N.D., (2004), "Reinforcement Learning for Reactive Power Control", *IEEE Transportations on Power Systems*, 19,3,1317-1325.
- Van de Wijdeven, J. and de Jager, B., (2005), "Shape Change of Tensegrity Structures: Design and Control", *Proceedings of the American Control Conference*, Protland, OR, USA, 2522-2527.

Waheed, A. and Adeli, H., (2005), "Case-based reasoning in steel bridge engineering",
Knowledge-Based Systems, 18, 37-46.

Tables

Table 1. Components of a case

Response of the structure	Control command	Slope compensation
$TSSD_{case}, v_{case}, TSD_{case},$	\mathbf{x}_{case}	$SC_m(\mathbf{x}_{case})$

Figures




Fig. 1. Five module active tensegrity structure: 15 m² plan area 10 active struts, 3 displacement sensors

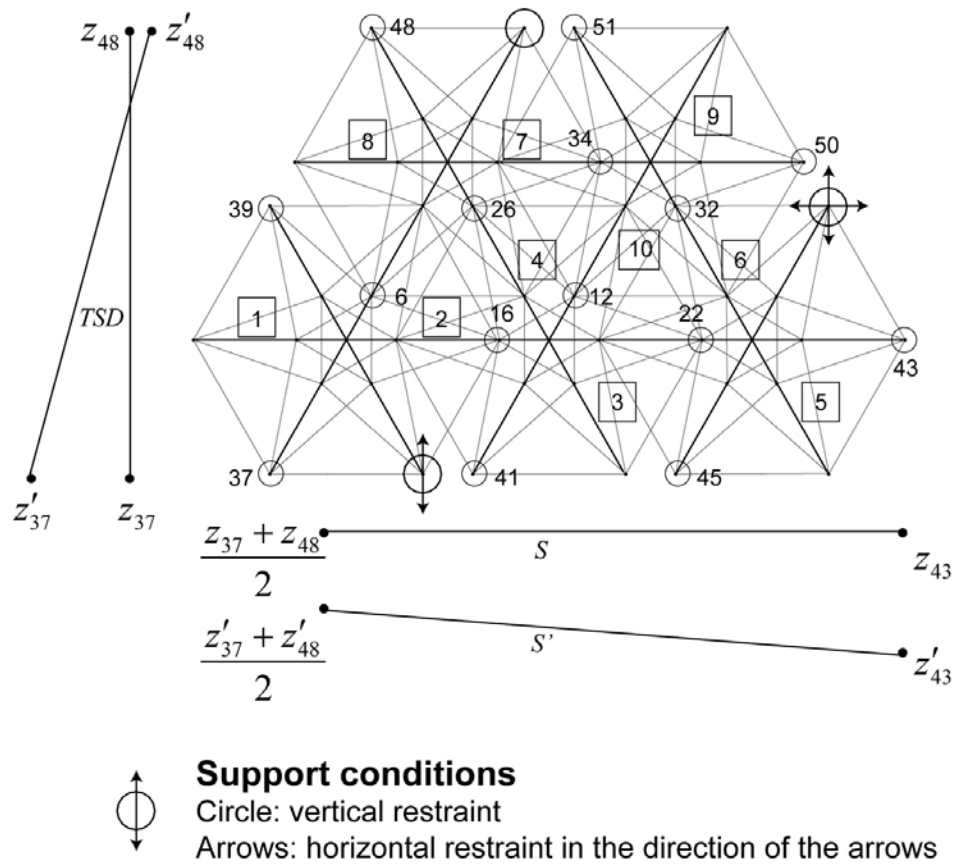


Fig. 2. View of the structure from the top showing slope indicators

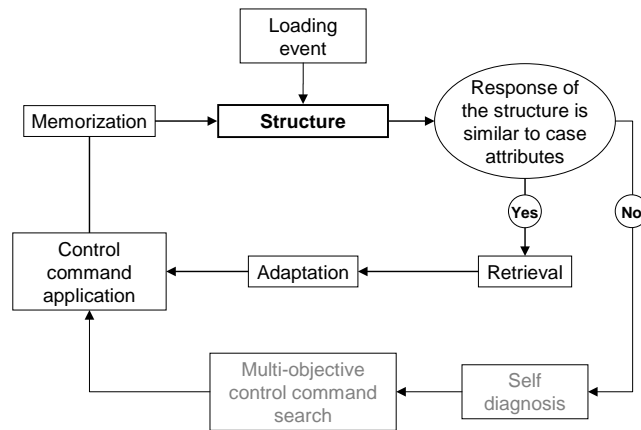


Fig. 3. Intelligent methodologies within the control loop




Fig. 4. Experimentally observed slope compensation improves over three different loading events that induce similar responses




Fig. 5. Error of slope compensation generally decreases over 80 loading and control events

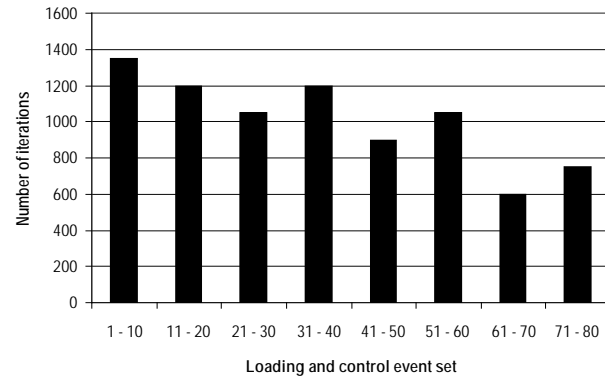


Fig. 6. The average number of iterations for 10 control commands generally decreases over 80 control events

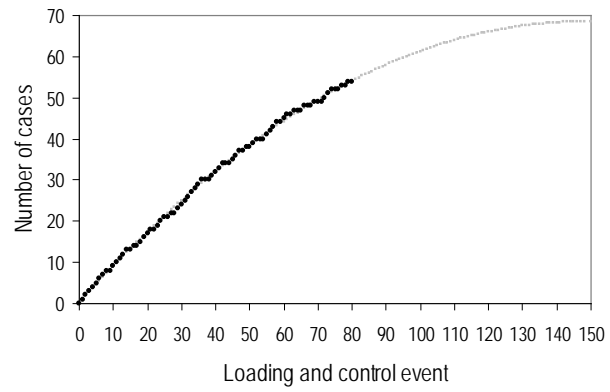


Fig. 7. Evolution of the number of cases over 80 loading and control events (black) and second order projection until saturation (gray)

Table captions

Table 1. Components of a case

Figure captions

Fig. 1. Five module active tensegrity structure: 15 m² plan area 10 active struts, 3 displacement sensors

Fig. 2. View of the structure from the top showing slope indicators

Fig. 3. Intelligent methodologies within the control loop

Fig. 4. Experimentally observed slope compensation improves over three different loading events that induce similar responses

Fig. 5. Error of slope compensation generally decreases over 80 loading and control events

Fig. 6. The average number of iterations for 10 control commands generally decreases over 80 control events

Fig. 7. Evolution of the number of cases over 80 loading and control events (black) and second order projection until saturation (gray)