

TOWSON UNIVERSITY
COLLEGE OF GRADUATE STUDIES AND RESEARCH

FUZZY MEDIATION AS A DYNAMIC METHOD OF INFORMATION FUSION
AND SHARED CONTROL

by
Giovanni Vincenti

A thesis
presented to the faculty of
Towson University
in partial fulfillment
of the requirements for the degree
Doctor of Science

August 2007

Towson University
Towson, Maryland 21252

© 2007 by Giovanni Vincenti

All Rights Reserved

TOWSON UNIVERSITY
COLLEGE OF GRADUATE STUDIES AND RESEARCH

THESIS APPROVAL PAGE

This is to certify that the thesis prepared by Giovanni Vincenti, entitled FUZZY MEDIATION AS A DYNAMIC METHOD OF INFORMATION FUSION AND SHARED CONTROL, has been approved by this committee as satisfactory completion of the requirement for the degree of Doctor of Science in Applied Information Technology in the department of Computer and Information Sciences.

Chair, Thesis Committee

Date

Print Name

Committee Member

Date

Print Name

Committee Member

Date

Print Name

Committee Member

Date

Print Name

Committee Member

Date

Print Name

Committee Member

Date

Print Name

Dean, College of Graduate Studies and Research

Date

ABSTRACT

FUZZY MEDIATION AS A DYNAMIC METHOD OF INFORMATION FUSION AND SHARED CONTROL

Giovanni Vincenti

Information fusion algorithms are designed to perform efficiently when dealing with multiple sensors and environments with a static set of integration rules. When placed in an environment with multiple controllers and the need for dynamic fusion, conventional algorithms do not extend to accomplish the task. Fuzzy Mediation is an innovative approach to creating an information fusion framework that can be used to mediate between an expert and a novice controller.

With this research, we explore different aspects of Fuzzy Mediation. We start with the demonstration of the concept in conceptual as well as static environments, and then we move on to simulations that place an expert controller and a one based on artificially intelligent algorithms at the commands of an agent. Finally we explore the reactions of human controllers to the Fuzzy Mediation framework in a setting of learning a new motor task.

The results show that Fuzzy Mediation is a successful extension of conventional information fusion techniques to accommodate for the dynamic interaction of two

controllers directing an agent. We observed that the overall performance of the agent based on Fuzzy Mediation shows an improvement when compared to the performance of an agent directed solely by the novice controller. Moreover, this framework allows for a novice with no training to start interacting with the agent after observing very few examples. This situation would not be possible if an unsupervised novice controller directed an agent. When human controllers interact, Fuzzy Mediation shows its strengths as an index of how closely the two operators perform.

The results obtained through this research show that Fuzzy Mediation is a very promising dynamic extension of information fusion techniques. As this is a new approach to shared control, it sets the perfect ground for further research. The possible applications range from the automotive and aeronautical fields to the evaluation of training effectiveness and performance.

Table of Contents

List of Tables.....	ix
List of Figures.....	xi
1. Introduction.....	1
2. Literature Review.....	6
2.1. Shared Control.....	6
2.2. Information Fusion.....	9
2.3. Applications of Information Fusion.....	10
2.4. An Introduction to Fuzzy Sets.....	14
2.4.1. Using Fuzzy Sets: an Example.....	18
2.5. Learning in Humans.....	21
2.5.1. Positive Reinforcement.....	21
2.5.2. Learning by Imitation.....	22
2.6. Summary.....	24
3. Methods.....	26
3.1. Fuzzy Mediation.....	30
3.1.1. Analysis of the Inputs.....	31
3.1.2. Revision of the Weight of Control.....	33
3.1.3. Calculation of the Single Output.....	34
3.2. Demonstration of Concept.....	35
3.3. Summary.....	40
4. Prerequisites.....	41
4.1. Simulation Environments.....	41
4.1.1. Simple Two-Controller Line Follower.....	42
4.1.2. Simple Flight Simulator.....	45
4.2. Learning in Agents Based on the Wang-Mendel Algorithm.....	51
4.2.1. Training Algorithms.....	53
4.2.2. Analysis of the Various Training Algorithms.....	54
4.3. Summary.....	59
5. Experiments and Results.....	60
5.1. Concept of Fuzzy Mediation Applied to a Two-Agent Control.....	63
5.1.1. Experiments with Controller N1.....	64
5.1.2. Experiments with Controllers N2 and N3.....	67
5.1.3. Experiments with Controllers N4 and N5.....	67
5.1.4. Recommendations.....	68
5.2. Fuzzy Mediation with Static FAM.....	68
5.3. Mediation Algorithm Experiments.....	73
5.4. Fuzzy Mediation with a Dynamic FAM.....	80
5.4.1. Preliminary Study.....	80
5.4.2. Observation of a Single Lap.....	83
5.4.3. Behavior of Different Granularity Agents.....	89

5.4.4. Behavior with Different Tightness of Control.....	94
5.4.5. Behavior with Different Mediation Algorithms.....	97
5.5. Fuzzy Mediation with a Simple Flight Simulator.....	102
5.5.1. Shared Control.....	103
5.5.2. Bi-manual Control.....	112
5.5.3. Observations.....	120
5.6. Interview with a Pilot of Submarines.....	120
5.7. Summary.....	122
6. Discussion.....	123
6.1. Possible Applications.....	124
6.1.1. Automotive technologies.....	124
6.1.2. Aeronautical Technologies.....	126
6.1.3. Medical.....	127
6.1.4. Training.....	128
6.1.5. Metric.....	129
6.2. Future work.....	130
7. Conclusions.....	132
Appendices.....	135
Appendix A – Institutional Review Board Documents.....	136
Appendix B – Responses to Pre- and Post-Study Questionnaires (Experiment 1).....	142
Appendix C – Responses to Pre- and Post-Study Questionnaires (Experiment 2).....	147
List of References.....	150
Curriculum Vitae.....	155

List of Tables

Table 1: Limits of fuzzy sets used for classifying the deviation between inputs.....	43
Table 2: Script for simulations studying shared control.....	46
Table 3: Script for simulations studying bi-manual control.....	46
Table 4: Average error for Equations WM1 to WM4, versus training algorithms.....	55
Table 5: Average error of Equation WM3 given different granularity and training.....	57
Table 6: Average error for Equation WM5, various granularities and training.....	58
Table 7: Equations used to power expert and novice controllers in the first experiment..	63
Table 8: Agent performance for one lap using different control levels.....	64
Table 9: Agent performance in section 2 of the simulation.....	66
Table 10: Boundaries of fuzzy sets used to classify the deviation between controllers....	70
Table 11: Core training for the sine environment with three different number of sets....	71
Table 12: Avg. mediator weight for 11-set agent (novice), different training algorithms..	73
Table 13: Agent performance over 10 laps.....	76
Table 14: Average and Standard Deviation values for the base experiment.....	76
Table 15: Performance of the M1 mediation engine.....	77
Table 16: Performance of the M2 mediation engine.....	77
Table 17: Performance of the M3 mediation engine.....	77
Table 18: Performance of the M4 mediation engine.....	77
Table 19: Performance of the M5 mediation engine.....	77
Table 20: Training pairs required for greedy training, different granularity agents.....	81
Table 21: Avg. abs. diff. between expected/actual output, different count training pairs..	82
Table 22: Avg. abs. difference between expected and actual output, different trainings..	82
Table 23: Evolution of the FAM as the simulation proceeds.....	85
Table 24: Final FAM.....	86
Table 25: Average absolute difference between inputs, Trainer/Trainee.....	90
Table 26: Standard deviations associated to values in Table 25.....	91
Table 27: Average absolute difference for output values, Trainer/Mediated Output.....	91
Table 28: Standard deviations associated to values in Table 27.....	92
Table 29: Average mediator weight.....	92
Table 30: Standard deviations associated to values in Table 29.....	93
Table 31: Average absolute difference between inputs, Trainer/Trainee.....	94
Table 32: Standard deviations associated to values in Table 31.....	95
Table 33: Average absolute difference between outputs, Trainer/Mediated Output.....	95
Table 34: Standard deviations associated to values in Table 33.....	96
Table 35: Average mediator weight.....	96
Table 36: Standard deviations associated to values in Table 35.....	97
Table 37: Average absolute difference between input, Trainer/Trainee.....	98
Table 38: Standard deviations associated to values in Table 37.....	98
Table 39: Average absolute difference between outputs, Trainer/Mediated Output.....	99
Table 40: Standard deviations associated to values in Table 39.....	100

Table 41: Average mediator weight.....	100
Table 42: Standard deviations associated to values in Table 41.....	101
Table 43: Performance of the participants by type.....	104
Table 44: Average performance of the experimental group.....	105
Table 45: Average performance of the control group.....	105
Table 46: Ideal run.....	107
Table 47: Performance of the control group.....	107
Table 48: Performance of the experimental group.....	107
Table 49: Performance of the experimental group during the shared control run.....	108
Table 50: Performance of the participants by type.....	113
Table 51: Average performance of the experimental group.....	114
Table 52: Average performance of the control group.....	114
Table 53: Ideal run.....	115
Table 54: Performance of the control subjects.....	115
Table 55: Performance of the experimental group.....	116
Table 56: Experimental group, run 3 (bi-manual control).....	116
Table 57: Interview with a pilot of submarines.....	121

List of Figures

Figure 1: Visual representation of a crisp set.....	15
Figure 2: Visual representation of a fuzzy set.....	15
Figure 3: Visual representation of a Triangular Membership Function.....	16
Figure 4: Visual representations of various membership functions.....	16
Figure 5: Sample calculation to find the membership value of a point.....	17
Figure 6: Visual representation of Type-1 (a) and Type-2 (b) fuzzy sets.....	18
Figure 7: Crisp sets representing categories of speed.....	19
Figure 8: Fuzzy sets representing categories of speed.....	20
Figure 9: Architecture for the training system.....	27
Figure 10: Fuzzy mediation system with two controllers.....	28
Figure 11: One-to-one training case.....	29
Figure 12: One-to-many training, parallel sessions case.....	29
Figure 13: One-to-many training, threaded sessions case.....	29
Figure 14: Fuzzy sets used for classifying the deviation between inputs.....	32
Figure 15: Fuzzy sets that regulate the balance of control.....	33
Figure 16: Simulation with two linear functions.....	36
Figure 17: Simulation with an expert linear function and a novice logarithmic one.....	37
Figure 18: Detail of the area in Figure 17 where control shifts to the novice.....	38
Figure 19: Detail of the area in Figure 17 where control shifts to the expert.....	39
Figure 20: Simulation with two sine waves.....	39
Figure 21: Detail of one of the areas in Figure 20 where control shifts to the expert.....	40
Figure 22: Path of the simulation.....	42
Figure 23: Diagram of the simulated agent with its sensors.....	43
Figure 24: Control functions, expert (dotted) and novice (solid) controllers.....	44
Figure 25: Experimental setup for the simple flight simulator.....	45
Figure 26: Diagram showing the simulation environment.....	46
Figure 27: Information flow architecture for the simulations.....	50
Figure 28: Shift in mediator weight during section 2, tight control.....	65
Figure 29: Shift in mediator weight during section 2, moderate control.....	65
Figure 30: Shift in mediator weight during section 2, loose control.....	65
Figure 31: Input and output values for a 7-set agent, core training.....	72
Figure 32: Control mediation value for 7-set agent, core training.....	72
Figure 33: Input and output values for a 11-set agent, core training.....	72
Figure 34: Control mediation value for 11-set agent, core training.....	72
Figure 35: Input and output values for a 21-set agent, core training.....	72
Figure 36: Control mediation value for 21-set agent, core training.....	72
Figure 37: Mediation engines.....	75
Figure 38: Single Lap Observation - Raw data between steps 1 and 100.....	87
Figure 39: Shift in control between steps 1 and 100.....	87
Figure 40: Raw data between steps 500 and 600.....	88

Figure 41: Shift in control between steps 500 and 600.....	89
Figure 42: Ideal path, distance from the center of the airfield.....	103
Figure 43: First run of test subject 6.....	108
Figure 44: Third run of test subject 6.....	108
Figure 45: Second run of test subject 6.....	109
Figure 46: Mediation weight $([-1, 1])$ and abs. difference between inputs (run 2).....	109
Figure 47: Performance of subject 4 in run 1.....	117
Figure 48: Performance of subject 4 in run 2.....	117
Figure 49: Performance of subject 4 in run 3.....	117
Figure 50: Performance of subject 4 in run 4.....	117
Figure 51: Mediation weight $([-1, 1])$ and abs. error between controllers (run 3).....	118

1. Introduction

Technology plays a dominant role as each new product is designed and placed on the market. We can find high-end technological solutions in any types of objects, from fuzzy logic-powered rice cookers to the advanced fly-by-wire controls in jet planes. If a technological solution oversees the cooking of rice, one may or may not require much improvement. However, when we talk about implementable solutions that control a passenger jetliner, we need to be sure that the concepts and solutions are well founded.

The fly-by-wire airplane idea is similar to the one of an automobile that operates with little input from the driver, or a joystick-operated heavy-duty machine that performs heavy jobs. What these three examples have in common is the existence of a computer mediating between signals from the operators and the actual sequence of actions taken by the machine in response to the operator. Sometimes, the response of the machine depends on the operator's input; however, there are times when the controls are autonomous responses to analyses of situations at hand.

Who (and when) should have control over the machine: the human or the automatic operator? Should we share control, and if so, how? Norman (2005) points out that what we understand as shared control is not really shared. For example, in early stages of development, the machine used to supervise and control any part of the operations in fly-by-wire airplanes, and now they oversee operations and still control flight, but as the conditions do not meet the standards of operations, the pilot is left alone. Therefore, it is

either the pilot controlling the plane on his/her own, or the control system, and there is no interaction between the two. We need a system that allows for greater interaction between the human operator and the digital one. There are examples (Caterpillar machines, for example) that allow operators to perform certain tasks only through automated systems (Grenoble O'Malley, 2005). Also, the two major airliner producers, Airbus and Boeing, are gearing towards advanced fly-by-wire technologies (Wallace, 2000), as the auto industry is attempting to infuse automation concepts in their products (Norman, 2005). These initiatives are important in the sense that automation is becoming predominantly part of everyday life. It is important, though, to realize that there needs to be some type of balance between controllers, let those be humans, computers, or a mix.

As the interaction between machines and digital controllers increases, we need to 1) find a better way to mediate control in dual control systems, when two (or more) operators are controlling the same machine, and 2) then, replace one of the human operators with a digital one and investigate the interaction between the two entities using a mediation system. This can be done if we: 3) create a framework for training using simulation and virtual reality to test-drive solutions and 4) implement these systems as a part of the actual operations of the machines under scrutiny.

This thesis introduces an innovative method that satisfies the needs just described, called Fuzzy Mediation. As we go on with a review of the literature on the subjects of shared control and information fusion the importance and distinctiveness of this algorithm will become extremely evident. As this is a solution that finds its niche in a

domain that has not been explored extensively, we will not be able to compare it to any other method already affirmed in the academic and industrial world. Instead, we hope that this work will lay the foundation to a long and successful stream of studies to be performed on concepts of dynamic information fusion coupled with issues in shared control.

The aims of the research are multiple. The first is the creation of a framework that is suitable for on-line learning. As we should always consider thoroughly all possible operating conditions when creating a new piece of software, also in this case we need to think of all possible uses. In this particular case, we envision the worst-case scenario as a situation where an expert controller has full knowledge and is trying to teach motor control to a novice that has no knowledge whatsoever. As we will see later in this work, we explored a few examples where the novice controller was completely unaware of any rules that governed the simulations.

Then, we want to construct an architecture that allows for true shared control. Some of the examples we will review later show that shared control resembles more an automatic control that allows for slight adjustments made by a second controller, usually human. This algorithm tries to break that barrier by introducing a dynamically balanced shared control.

This also identifies another goal of this work, which is to extend the field of information fusion to host a set of algorithms that are dynamic in nature. Some evidence of the lack of truly dynamic information fusion will be given in the next chapter, but it is important to note that not all environments remain static in operational requirements. As

the necessities change, also the algorithm and the mediation of control should adapt, keeping the object's operations safe.

The last aim of this research is to use Fuzzy Mediation as a new metric to assess the closeness of control in several environments. More on these possible applications of this algorithm will be introduced in the sixth chapter.

This thesis will continue in the following manner. The next chapter is dedicated to the review of concepts that build the foundations of information fusion. We will analyze in particular the domains of shared control and information fusion. We will also review some applications of information fusion to show the importance of this field.

In the following chapter we will introduce the concept of Fuzzy Mediation in great detail, also showing a proof of concept that should make the reader at ease with both the technical articulations as well as the overall solution that this algorithm offers.

In the fourth chapter we review some information that needs to be addressed before we can continue with the analysis of the experiments performed. Especially, we will take a look at the simulation environments created for this particular purpose, and we will also review the operations of an agent based on the Wang-Mendel algorithm.

In the fifth chapter we will review the results of several experiments performed using Fuzzy Mediation as an information fusion engine. We will explore solutions with static, dynamic and human agents that try to navigate through the simulation environments as well as some components that affect the success of this mediation algorithm.

In the following chapter we will discuss potential practical applications of this work,

which will extract the reader from the technicalities associated with these writings, and will project Fuzzy Mediation in a domain that can truly make full use of its potential.

Finally, the last chapter will conclude this thesis.

2. Literature Review

Before we get any further into the analysis of the methods involved with Fuzzy Mediation, it is important to review some of the concepts that were the basis for this algorithm to grow.

First, we will look at different approaches to shared control, which constitutes the major influence for the potential applications of this work. Then we will look at the field of information fusion, responsible for the integration of several inputs to generate a single and coherent output. We will also review some of the applications where information fusion plays a dominant role. Finally, given the nature of the method proposed in this work, we will review briefly concepts of fuzzy set theory and its practical applications.

2.1. Shared Control

Shared control is a domain of research that spans over many different kinds of applications. In this section we will review a few approaches to the solution of this problem. It is important to define immediately that shared control is based on two key points, as defined by Wasson and Gunderson (2001). The first concept is that there needs to be some degree of communication between the two controllers, the second is that there should be a way to understand the goals of the other agents involved in the system. Wasson and Gunderson (2001) also bring up the point that, when the controllers are automated agents, it is simple to create some type of communications standards that all agents need to conform to and where they need to rely any information that they have. When the agents are both human and automated, problems arise.

The first work that we will analyze is the one by Wasson et al. (2001, 2003), where they devised a set of algorithms that allow for the elderly to maneuver efficiently and in collaboration with a computer system a personal mobility device, or walker. In such system, the authors state that, given the nature of the walker and the instability of the patient, it is of crucial importance for the human to get a sense of shared control (Wasson et al., 2001). The authors also state that the changes perceived by the autonomous controller must be incorporated into the overall output to allow for the person not to be disoriented by sudden changes in direction of the walker, thus provoking a loss of balance. The type of collaboration that Wasson et al. (2003) refer to then is simply in the fine adjustments of the controller's final output, limiting greatly the collaborative aspect of control.

The research by Wasson et al. (2003) brings out several major control modes that the controller must follow, and they are the warning system, the safety braking, the safety braking and steering, and the path following modes. In the warning mode, the walker can only alert the user as to problems that may impede a smooth walk. In the safety braking, the walker will detect obstacles and simply break when the danger becomes too grave to the balance or when the person is stopped. The third mode, the safety breaking and steering comes into play when the walker needs to respond to a potentially dangerous situation by applying both braking as well as steering maneuvers to keep the person safe. Finally, the path-following mode works as the safety braking and steering, except it is on at all times, not only when a danger is sensed.

The advantages to this algorithm are many. First of all, the user feels in control, even though rigid mathematical equations are in place to maintain the user along a

programmed path. Also, the smooth collaboration allows for the user's balance to stay somewhat intact as the walker performs its functions. The most apparent drawback that shows from analyzing the algorithm is that the deviation between the user's control and the walker's control must be of zero for the user to really be in control.

Other projects that built on this concept include the later work of Wasson and Gunderson (2001) and Wasson et al. (2003). In the work by Wasson and Gunderson (2001), it is defined that the walker has four states, "No Assist", "Assist", "Safety" and "Override". These four states can easily be traced back to the four control modes defined by earlier research (Wasson et al., 2003). Especially in this case, it is apparent that the transition between states is immediate, and the state of "No Assist" is reached only in cases where there are no obstacles that can be sensed. This suggests that, any time some danger situation is present, the algorithm leaves little to no trust to the user, thus overriding "gently" the majority of the actions.

Other projects in similar settings include the works of Gomi and Griffith (1998), Levine et al (1999) and Yanko (1998), with the creation of "smart" wheelchairs. Other work along similar lines was created by Connell and Viola (1990), who realized a semi-autonomous mobile robot based on collaborative control. A different version of collaborative control stemmed into the realization of the GuideCane by Borenstein and Ulrich (1997), a special cane that aims to helping blind pedestrians in walking. An example of a more elaborate project that involves constant path re-planning as well as self-localization is the Flo project (Roy et al, 2000), a robotic nurse that collects data about patients and gives minimal robotic assistance, without providing physical support.

Steering away from the field of shared control applied to the elderly, we find many other projects that deal with shared control. Perhaps the most evident need for control shared between a human and an automated controller comes when heavy material needs to be handled. Colgate, Wannasuphoprasit and Peshkin (1996) touch on this scenario by means of their cobots, robots meant for collaboration with human operators. Also the cobot approach is heavily relying on strict mathematical association between the human controller and the artificial one, still leaving only adjustments for humans to perform along the automated operations. An example of a heavy-material handling robot is described by Shinohara, Miyawaki and Sunayama (2001).

2.2. Information Fusion

Control is a concept that involves the interaction of multiple entities, by definition at least two. In any situation of control we find at least one subject who interacts directly with one object, directing the second object's every move (WordNET, n.a.). In environments where there is more than one operator controlling one object, we turn to the discipline of information fusion for solutions on joint control. Information fusion is defined as the process of taking multiple inputs and creating a single output (Kokar et al., 2004). As Kokar states, many information fusion algorithms are biased in their operations. As they are designed, their operations are set and the fusion is carried out by simply running the algorithm. In most cases, the algorithms are pre-defined and static, and there is no adaptivity to the circumstances that influence the object. Adaptive information fusion systems, on the other hand, would offer an alternative where control can be effectively and truly shared between two (or more) operators.

A shared-control environment would allow for higher-level applications to higher levels, in cases where supervised and collaborative learning is required, as in Trajkovski (2005). Supervised machine learning techniques represent a set of operations that "learn a task starting from a suite of examples" (Abad et al., 2002). Information fusion would provide for supervised machine learning algorithms that learn from examples as the expert (operator) is operating (controlling) the machine. This approach represents a paradigmatical shift, as the learner becomes capable of interacting efficiently with the environment. It is an effective shift from the supervised learning domain to the collaborative learning one (Gokhale, 1995; Hammonds et al., 1997). It has been noted, that many proposed solutions promote individual, instead of true collaborative learning (Dugan & Glinert, 2002). Kwek (1999) explored interaction between humans and computing when learning is performed by means of an apprentice model; however their focus is on the classification of items.

2.3. Applications of Information Fusion

Information Fusion is a discipline that lends itself to many applications. In this section we will analyze some of them with the purpose of showing how this discipline adapts itself to many domains.

More than the simple application of information fusion techniques to real-life problems and solutions, we will also review the pivotal role that these algorithms play in the life of the new systems, given the key element of taking information from multiple sources and creating a single output.

The first application of Information Fusion that we will explore is the solution proposed by Ross and Jain (2003). These authors investigated the use of Information Fusion techniques applied to biometric devices. These kinds of devices are the ones that authenticate a person's identity by detecting physical characteristics that are unique to an individual (Shelley et al., 2007).

These devices are generally used as a way to lock computer functions from the use of the general public, or to maintain strict access policies to rooms, to name a few (Liu & Silverman, 2001). Examples of biometric scanners include fingerprint scanners, hand-geometry scanners and iris-recognition software (Shelley et al., 2007). Each one of these elements detects the singular characteristics of a person, such as a fingerprint, and looks for a possible match among the data stored in its databases. If the identity of the person is stored, and there is a match with the reading, then the identity is verified. Instead, if there is either a mismatch between who the person says s/he is and the identity stored in the database, or if the database does not contain any information at all, then the identity is not confirmed, and the system doesn't let the person perform the desired operation.

Typically, only one of these biometric devices is used to maintain security, and generally it is enough. Ross and Jain (2003) advance the theory that, should one specific type of biometric recognition be bypassed, the system would be open to intruders. For this reason, they propose a solution that utilizes Information Fusion techniques to merge information from several biometric recognition devices in an effort to make the identification system more secure and less prone to spoofing.

In their work, Ross and Jain (2003) joined three different kinds of biometric

authentication – face, fingerprint and hand geometry – in order to create a higher level of security. They show that using a simple face geometry algorithm alone may lead to problems, since face recognition systems may be affected by different lighting conditions. Also, partial fingerprints may allow for fingerprint recognition not to work as expected. The reading of three different biological aspects of a person was shown to be much more effective in the identification process, even if some characteristics may not lead to a perfect match, through Information Fusion techniques.

The second example in the application of Information Fusion techniques is the one advanced by Zhang et al. (2005), who proposed the use of such techniques in combination to Genetics Programming in order to improve the classification of text documents into predefined categories.

The first important distinction to make is that the authors use Genetic Programming rather than Genetic Algorithms. When using genetic algorithms, the user of the system is limited to simple data to represent the genetic makeup of an object (Ribeiro Filho et al., 1994). Instead, with genetic programming, the system can utilize more complex data structures, such as trees, lists or stacks (Zhang et al., 2005).

This research is pushed by the fact that, the majority of the times, just reading the title or the abstract alone of a document is not sufficient to categorize appropriately the type of content presented. Also, generic voting systems that power current classification engines seem not to work well enough for the classification of text documents. The core idea that powers the solution created by Zhang et al. (2005) utilizes each section of the text document as a different object. Once all the documents have been cataloged and

stored into memory, then the genetic programming component of the system calculates the similarity value to each section of the document. Since this system is taking into consideration more than one element of the document, Information Fusion offers the perfect framework to analyzing and creating a single similarity value from several different components of the document. This method of classification allowed for a more appropriate classification of the documents used.

The last example that we will explore is the use of Information Fusion techniques to the precision guidance of agricultural vehicles, as reported by Reid (1998) and Rovira Màs et al. (2005). Agricultural vehicles operate in a condition for which it is not possible to use lines on the asphalt or GPS alone to allow an automated controller to drive. As many components play a role in the correct navigation as well as operations of an agricultural vehicle, then Information Fusion offers a great system to take in several guidance mechanisms and create a single output that will keep the vehicle operational.

Automated guidance systems typically use a set of different controllers that interact and allow the correct operation of the vehicle. Such controllers can perform the following functions: path planning, control of the steering, posture sensors and feedback controllers. A more complex framework should really be in place, where several other mechanisms are collaborating. Reid (1998) revisits an architecture by Noguchi et al. (1998) where the fusion allows inputs from a geomagnetic direction sensor, a GPS system and a machine vision system to collaborate to accomplish the task.

Reid (1998) and Rovira Màs et al. (2005) highlight that the future of automated

guidance of agricultural systems is very interesting with the development of several new mechanisms that would allow vehicles to become autonomous. On the other hand, it also poses a very tough situation with an increasing need for the creation of platforms that can perform the fusion of information from several different controlling devices, in order to obtain a better guidance system.

As these are just three examples of applications of Information Fusion, this section wants in no way be an exhaustive review of such topic. What is important though is the fact that fusion techniques have very practical applications that not only perform some function, but the overall result is extremely practical in its application.

Also, we reviewed applications in biometric identification, text classification and the guidance of an agricultural vehicle. These examples provide the wide range of acceptance of Information Fusion techniques as a viable solution to the problem of merging several different controlling signals and creating a single output that is precise and clear.

2.4. An Introduction to Fuzzy Sets

When we think of mathematics, we think of a precise science that usually has a one-number answer to most questions. Fuzzy sets fit this stereotype as well, even though the name is somewhat misleading in this regard. Fuzzy sets belong to a branch of mathematics that deals primarily with set theory (Pedrycz & Gomide, 1998).

Every number can be considered a part of a set. The common understanding of classic set theory is that a value either belongs to one set or it does not. Such sets are referred to as crisp sets (Pedrycz & Gomide, 1998). A visual representation of this

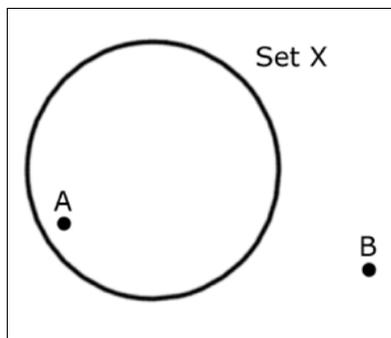


Figure 1: Visual representation of a crisp set

concept is reported in Figure 1, which shows that A is part of Set X, and B is not.

In fuzzy sets, each value can have a partial membership to one set (Pedrycz & Gomide, 1998). This property allows us to create sets where numbers can refer to one set or another, since there is no limit to how many sets a value can have a partial membership to. Figure 2 shows how Figure 1 could be modified, if set X were to be a fuzzy set partially including A. Even though B still does not belong to set X and A still belongs to it, this time point A does not belong to it completely. The shaded area defines the rate at which a set includes the values in that area. Since point A is part of that shaded area, and not the fully dark area (center of the set, where every item has a membership of 1 to that set), its membership value to set X is partial, say 0.75. Every fuzzy set can be represented in this way, and every area where values belong to one set only partially can be represented by functions that evaluate the membership coefficient of one value to the set

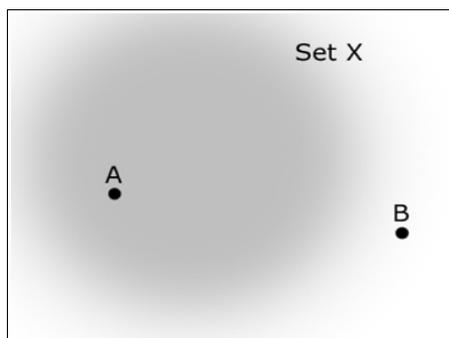


Figure 2: Visual representation of a fuzzy set

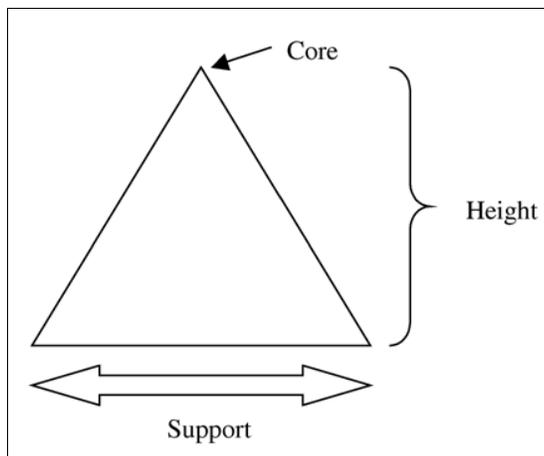


Figure 3: Visual representation of a Triangular Membership Function

under scrutiny. Every fuzzy set of the type discussed so far can be represented by a two-dimensional graph. Figure 3 shows the visual representation of a Triangular Membership Function. A range identified solely by sets with triangular membership functions is called a triangular partition with evenly spaced midpoints, or TPE (Sudkamp & Hammell, 1994). Such function is just one of various types of membership functions (Yen & Langari, 1999). Other types of membership functions are, but not limited to, bell-shaped, Gaussian, trapezoidal and sigmoidal, as shown in Figure 4 (Yen & Langari, 1999).

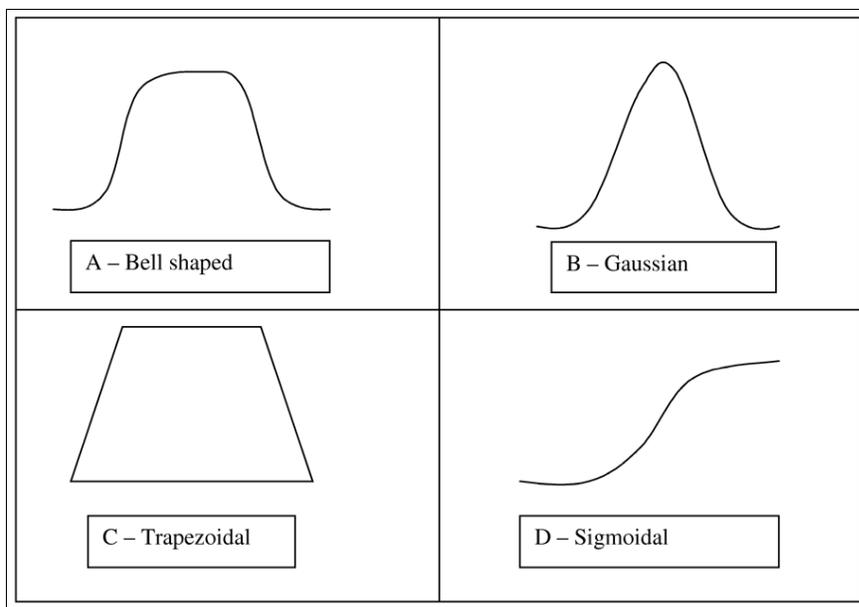


Figure 4: Visual representations of various membership functions

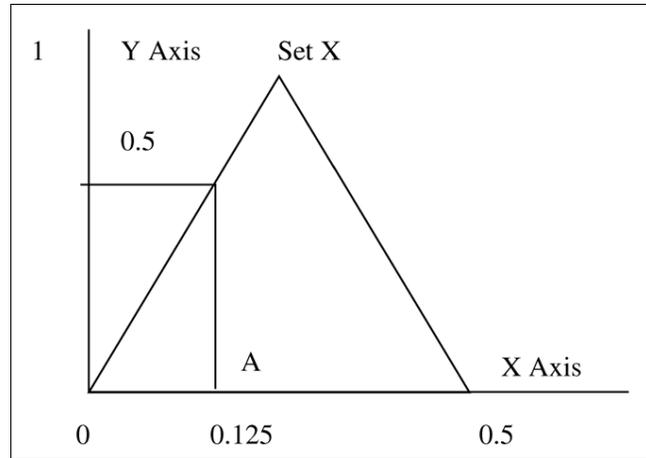


Figure 5: Sample calculation to find the membership value of a point

In order to calculate the membership value that each value may have to one fuzzy set, we have to find the projection on the y axis of the value that is included within the set along the x axis, as it is shown in Figure 5. We can assume that we are trying to find the membership value of point A to set X. We can also assume that $A = 0.125$ and the support for $X = [0, 0.5]$. Given this information, we find the value of A along the x-axis. As we find this spot, we will find the coordinate of the edge of the fuzzy set with the same x value. When we find such point, the y-value composing the coordinate will give us the membership value. In this case, point A's projection on the Y-axis given a Triangular Membership Function is point 0.5. This tells us that point A's membership value to set X is 0.5. We can use this system to find the membership function of any point, given the support of a set and the type of membership function that the set may show.

When we deal with fuzzy sets, we can have several kinds of membership functions, as we saw earlier. Those kinds of functions can be defined of type-1. Fuzzy sets that utilize type-1 membership functions are classified as Type-1 fuzzy sets. For this simulation we use these kinds of fuzzy sets. Opposed to type-1 membership functions, we

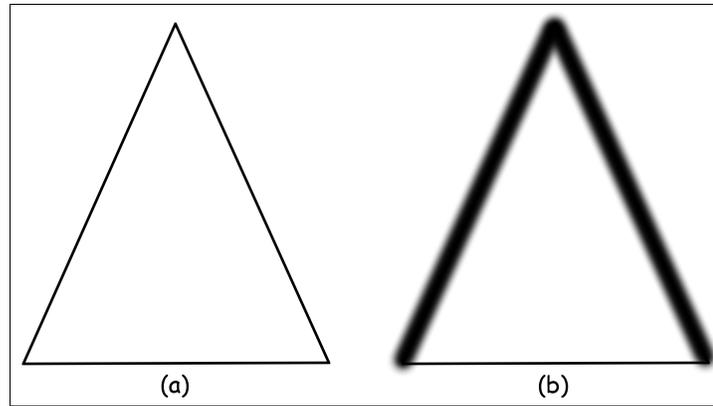


Figure 6: Visual representation of Type-1 (a) and Type-2 (b) fuzzy sets

find functions of type-2 (Mendel and John, 2002; Mendel, 2003). Such functions are more complex in nature. If type-1 functions are very clear, type-2 ones add a level of complexity. Figure 6 shows a representation of the difference between these different types of membership functions.

Type-2 membership functions do not operate in 2-dimensions, such as the ones of the first type, but they use a third dimension. Although researchers argued that Type-2 fuzzy sets are more appropriate in some situations where uncertainty is more predominant (Mendel and John, 2002; Mendel, 2003), we feel that the setting of these experiments is simple enough where the uncertainty of the boundaries of the fuzzy sets can easily be overcome by increasing the number of fuzzy sets associated with the input or the output ranges in question.

2.4.1. Using Fuzzy Sets: an Example

Now that we have at least a vague concept about fuzzy sets, we can analyze a problem and place it within a fuzzy set frame of reference. For this example, we can assume that we are trying to categorize the speed of a car into three sets: Slow Speed, Medium Speed, High Speed. If we were to represent such sets within a crisp

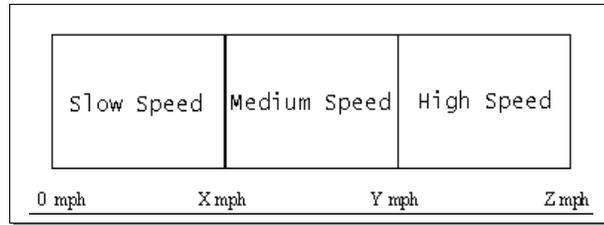


Figure 7: Crisp sets representing categories of speed

environment, Figure 7 gives us a feeling of what these sets may look like. As soon as we try to represent the concept of different levels of speed, we run into problems. As we can see, Figure 7 gives very clear boundaries for each of the three sets.

Every person will have a different interpretation of what speeds the “Slow Speed” set will include though. For example, some people may say that you can consider a speed to be slow when it is anywhere between 0 and 20 miles per hour. If asked to another person, the set of slow speeds may include any value between 0 and 30 miles per hour. So a speed of 15 mph will be considered by the majority to be a slow speed. Likewise, a speed of 40 mph will be considered a medium speed. We have problems in classifying the values that fall within areas of uncertainty, given to divergent opinions to what a set should include. For example, some people will see 25 mph as slow speed, and others will classify it as medium speed.

Fuzzy sets help us with this issue, as they take into consideration both sets when it comes to the classification of the value 25 mph. In fuzzy sets, if a number belongs to a set completely it is considered to have a membership of 1 to that set, instead, if a value is not included at all into one set, it is considered to have a membership of 0 to that set. The values that we have analyzed earlier are considered to belong to one set only, as there is no difference of opinions to the range of the small and medium speed sets. Each of these

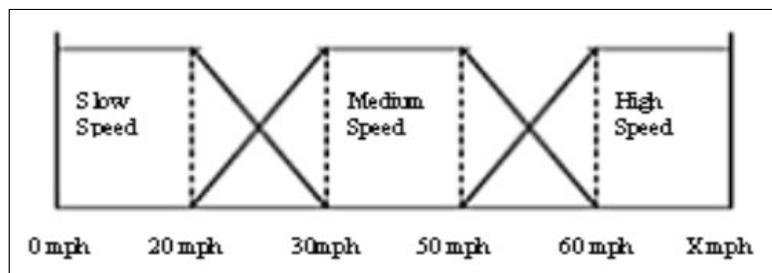


Figure 8: Fuzzy sets representing categories of speed

values will belong solely to one set, and we can say that the value of 15 mph has a membership value of 1 to the “Slow Speed” set, and a membership of 0 for the “Medium Speed” and “High Speed” sets. Likewise, the speed of 40 mph has a membership value of 1 to the “Medium Speed” set, and a membership of 0 to the “Slow Speed” and “High Speed” sets.

In order to understand this concept better, we should analyze the value of 25 mph. We can see from Figure 8 that the value falls in between the lower limit of “Slow Speed”, represented by the lowest value of the high end of the “Slow Speed” set (the first person thought that slow speeds ended at 20 mph) and the higher limit of this set, represented by the highest value for the high end of this set (the other person thought that the slow speed ended at 30 mph). Being the value of 25 mph right in the middle, we can say that it belongs to the “Slow Speed” set some, and to the “Medium Speed” set some as well, in equal parts. As we stated earlier, the maximum membership value that a number may have when considering its membership to a fuzzy set is 1, and if we consider a speed of 25 mph in the middle of the transition area between a “Slow Speed” and a “Medium Speed” set, we can say that the speed of 25 mph has a membership value of 0.5 to the “Slow Speed” set, and a membership value of 0.5 to the “Medium Speed” set.

This example is based on the assumption that the fuzzy sets have trapezoidal

membership functions, that the maximum membership value that a number can have to all of the sets it belongs to is 1, and that the value can belong to zero, one or two sets.

2.5. Learning in Humans

Nowadays no one argues the fact that the connectivity of neurons could not be specified in the human genome. It is widely believed and accepted that learning in the brain resides in the plasticity associated with alterations in synaptic efficacy.

2.5.1. Positive Reinforcement

The most widely accepted theory on the matter of learning at the cellular level is aligned with Hebb's reinforcement ideas (1949). This theory introduces the possibility that, as neurons fire, the synaptic connection between them is strengthened. The stronger synapses in turn lead to preferred pathways, which seem to then create memories. This process is called long-term potentiation (LTP) and has been widely investigated in Kandel's work, reported in Kandel et al. (2000). Given the nature of this topic, much mathematical modeling has been drafted. Among the many, neural networks have been widely accepted as the closest representation of the concept of LTP. Among neural networks, we should mention the "attractor" neural networks, or ANNs, and various forms of feed-forward networks trained by back-propagation algorithms. ANNs were introduced by John Hopfield (1982) as a response to the then current faulty models that did not represent well the physiological basis of learning, especially true with back-propagation networks. This is because these models don't explain the interaction with another brain, and the inherent lack of self-organization.

In reinforcement learning, the environment acts as a critic rather than a teacher. The downside of this approach though is that the learning is slow and it is not portable to a new task, so the subject will have to start all over. Chialvo and Bak (1999) argue that this is true because of the lack of positive reinforcement rules. They say that these models thrive on the concept of LTP, but they also state that "long-term synaptic depression (LTD) in the mammalian brain is almost as prevalent as potentiation, but there appears to be little or no understanding of its functional role." Barnes et al. (1994) argue that "although it is conceivable that LTP is the critical phenomenon used for storing information, and that LTD may exist simply to reset LTP, it must be noted that it is also conceivable for the converse to be true." Thus, Chialvo and Bak infer that the "depression" of synaptic efficacy is the fundamental dynamic mechanism in learning and adaptation, with LTP playing a secondary role. The collaboration of these two learning mechanisms produces the outcome that we are all familiar with, the acquisition of a new task. It is important though that LTP and LTD collaborate closely though, because it is much more likely that, as we learn, we perform some actions that are not successful, thus the need to "erase" that action from the set of successful ones through LTD, stimulating only the pathways that have given good results through LTP.

2.5.2. Learning by Imitation

The phenomenon of imitation was never seriously considered by scientists until the beginning of the 21st century. It was mainly considered to be an unintelligent process in higher primates, and, as such, not worth researching. After its earliest consideration by Thorndyke (1898), it only appears sporadically in psychological literature. The only

notable consideration was done by Piaget (1945) in his consideration about the developmental stages in children. Thorndyke considered imitation to be learning by observation, when one entity tries to mimic (copy) another entity's behavior.

With the discovery of mirror neurons by Rizzolatti et al. (1996), it seems like, after all, we are wired for imitation. The mirror neurons are located in Broca's F5 region in the frontal cortex, that has been found to be primary responsible for human linguistic expression. People with defects in Broca's region are usually not linguistically competent. In the observations of mirror neurons in primates, it has been noted that they fire in a monkey that is observing another monkey tearing or crumbling paper.

With the discovery of mirror neurons and the possible extent of their significance, imitation becomes a bona fide focus of research in learning. The indication for neonatal research that it is by imitation that we start learning everything about the world and build our first conceptualization of it. Language comes later in the process.

Therefore, it seems that imitation could be the base of all learning that happens in the human and it is connected to our very basic biological self. As we know, in learning the more basic the motivation is, the easier it is to assimilate new knowledge towards the satisfaction of this drive (Trajkovski, 2007). The main thesis of that work is that humans learn about the environment via interactions with it based on their inborn schemes in Piagetian sense, and from other humans via imitation conventions, when considered parts of a homogeneous multiagent society.

Mirror neurons as hardware for imitation can explain phenomena such as empathy in humans. Rizzolatti et al. (1996) discovered that a specific area of the brain, called F5, is involved with imitation of tasks. The neurons in this area are extremely active in

macaque monkeys when they are performing movements that are goal-oriented. Such neurons were observed to be extremely active also when the monkey observed another monkey or a human perform those very movements. This finding, together with Kandel's long term potentiation and the strengthening of synapses through repetitive stimulation, hints to the fact that observation and imitation are phenomena that lead to learning a task.

This past decade has been a great time for research in imitation. Heyes (2001) reports that advancements were made under various biological aspects. From an evolutionary standpoint, chimpanzees and birds were successfully studied for imitation patterns. Moving onto humans, the first point of view analyzed was the developmental one, where newborns were studied for the reflex of sticking out their tongue when an adult performs that action in front of them. Further studies were conducted on 18-month-old infants to study their imitation of movements that adults reported being intentional. When asked, the infants could not always motivate the performance of the imitated movements. Finally, Heyes reports that studies have shown that autism is linked to deficits the subjects have in the processes of imitating the performance of tasks they observe. Other studies were also performed on adults as well as at the organ level, with the study of the brain and patterns associated with task performance. The approach that we are presenting in this paper is based on learning by imitation in adults.

2.6. Summary

In this chapter we reviewed several concepts that lay the foundations for the ideas discussed next. The first topic discussed foundations and examples of shared control, which is perhaps the keystone of Fuzzy Mediation. Then we reviewed elements of

information fusion. This branch of computing allows us to take several inputs and generate a single, coherent output that we can use. These concepts lay the foundation for the core of Fuzzy Mediation. In order to put the practicality of information fusion in perspective, we also reviewed some practical applications of these concepts. Finally, we reviewed the basics of fuzzy set theory. Also this time, we explored a possible application of such theories in order to review the practicality of the concepts.

In the next chapter we will introduce the concept of Fuzzy Mediation in a very concrete way, exploring first the high-level concept and need, and then focusing on the more technical aspects of this approach.

3. Methods

This section introduces the ideas behind Fuzzy Mediation and gives an overview of the settings that will benefit from this approach to information fusion. This concept is the brainchild of the author of this work, and, to our knowledge, it is not an extension of any existing algorithm.

It is important to understand that an automated controller rarely performs better than the most complex situation envisioned by its creators, while humans astonish us every day with resourceful ways to solve the most complicated situations. Most collaborative environments allow for humans to apply only minor changes to the operations of an automated controller, placing de facto a machine's preprogrammed operations over the judgment of a human. Thus, there is the need for a new platform that allows for a more liberal collaboration between controllers of different nature and expertise trying to manipulate the same object.

The problem of shared control is most relevant in a training setting. This research aims to create an efficient training framework for operating machinery using a computer simulation (virtual environment). Initially we will focus on training car drivers, using a positive reinforcement approach, in a trainer-trainee shared control collaborative environment. This collaboration will be achieved greatly by devising the tool based on our original Fuzzy Mediation approach. The setup involves two controllers that are trying to steer a car within a virtual environment. One of the controllers is an expert driver, and the second a novice. When both subjects are trying to control the car, whom should

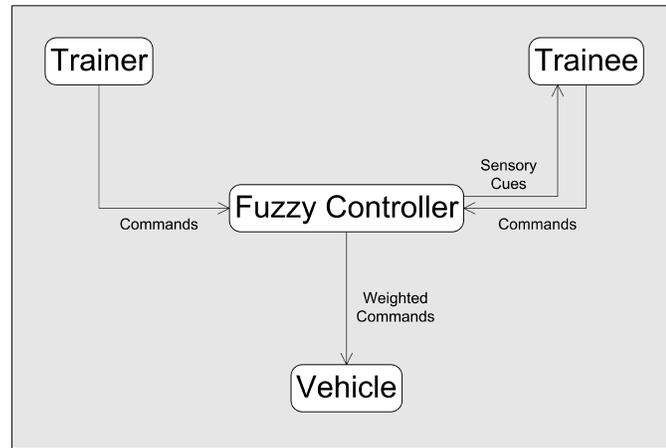


Figure 9: Architecture for the training system

control go to? If the expert drives, the novice will passively observe the expert's driving, and not really learning how to drive – the novice is just being exposed to examples of driving situation cases. If the novice driver is (fully) in command, the expert can only interact with the novice by simply correcting his/her actions by negative reinforcement.

Our proposed approach shares the control between the expert and the novice by stimulating the involvement of the novice. Using a Fuzzy Mediation control, inputs from both participants in the process are evaluated and considered, and the overall control gives more or less weight to the novice based on his/her past performance. This approach is a generalization of the negative reinforcement approach – the classical approach is just a marginal case of this one. The overview of the approach follows below.

This solution, depicted in Figure 9, involves both users at the same time. The expert and the novice are given a set of identical controls, and they both try to control the very same car within the same simulation. The inputs received by the controls of the car are analyzed and compared to the other participant's (novice to expert, expert to novice). The concept of the "Fuzzy Controller" will be explained in detail in Section 3.1. Initially, the

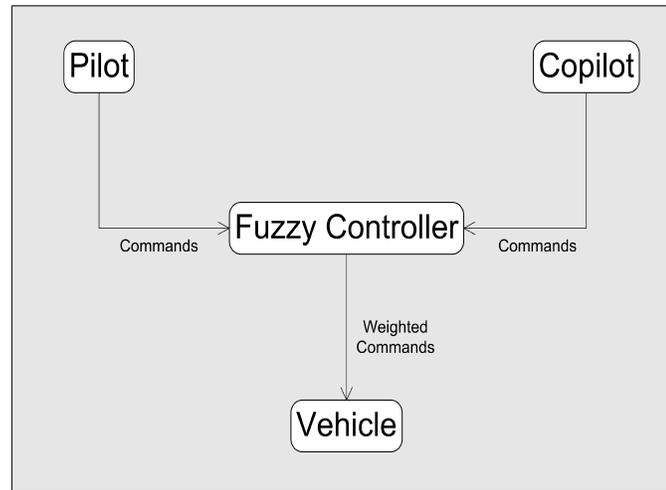


Figure 10: Fuzzy mediation system with two controllers

mediation system allows the vehicle to be controlled by the expert. As the novice's actions mimic/imitate more and more the correct ones of the expert, control shifts to the novice. Visual clues (indicators) are given for the novice to indicate progress (i.e. shifting control towards his/her side). Should the novice's actions start deviating from the expert's, control shifts back to the expert. As previously discussed, based on results on general success in positive reinforcement learning (Chialvo & Bak, 1999), we expect the novice to learn significantly better than in a negative reinforcement environment.

An extension of the system architecture in Figure 9 is given in Figure 10, when the control mediates between a digital and a human controller, suitable for, for example, some fly-by-wire airplanes in service today. In this case, the pilot is the human controller, and the copilot can be either the human co-pilot or a digital smart box.

This work is also adaptable to alternative architectures, shown in Figures 11 through 13, in environments with one expert and multiple novices are being trained. It is important to keep in mind different training architectures, given different types of learners.

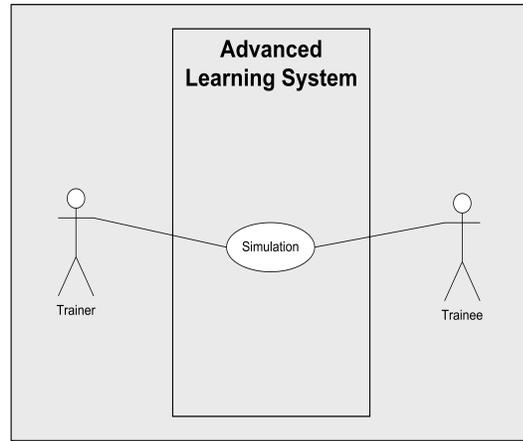


Figure 11: One-to-one training case

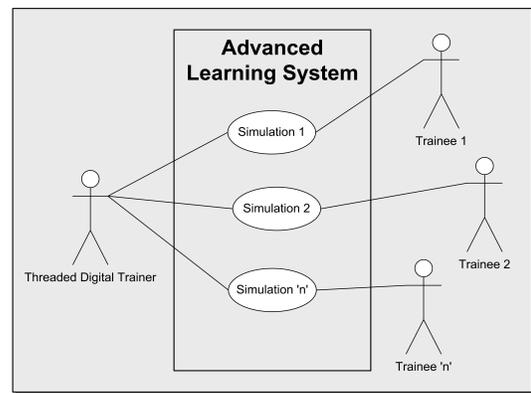
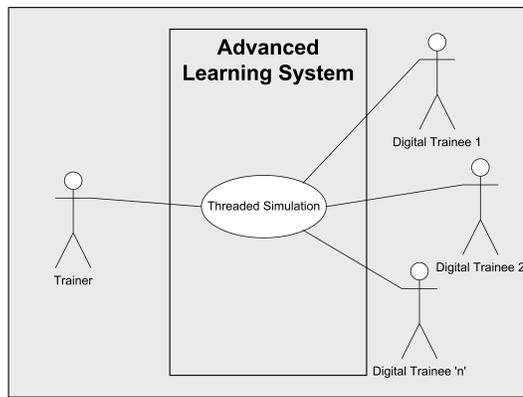


Figure 12: One-to-many training, parallel sessions case Figure 13: One-to-many training, threaded sessions case

In the most traditional case, learning happens on a one-on-one basis. This is the basis for the architecture shown in Figure 11. This type of environment is ideal for humans, where one expert and one novice interact. For this work we will focus entirely on this architecture.

Then we have to take into consideration different architectures where there is more than one novice at a time. Figure 12 introduces the first of the two architectures dedicated to multiple sessions at once. In this particular architecture we can see that a single expert interacts with a single simulation, where multiple novices are connected. This situation is

ideal for a human expert that is interacting with a single simulation, with either multiple humans or multiple A.I.-based novices. The most relevant point of this architecture is that there is a single simulation environment, and all novices have to adapt to the same motions of the expert.

Finally, Figure 13 shows the last of the three architectures proposed. In this architecture we can see again a single expert and multiple novices. In this particular environment though there is a dedicated simulation for each novice. This environment is ideal for either human or A.I.-based novices, with a single expert. In this second environment the expert would need to keep track of several simulations, not necessarily all at the same state. For this reason, the expert for this architecture is most likely A.I.-based.

3.1. Fuzzy Mediation

Information fusion is a domain that deals with the integration of input from multiple sources in order to create a single output (Kokar et al., 2004). This task is particularly difficult in the event of conflicting information coming from different sources. The mediation that needs to take place is the heart of any information fusion mechanism.

One of the applications that may come to mind when talking about information fusion systems is the domain of shared control. As technology advances, more vehicles rely heavily on computing for their control (Grenoble O'Malley, 2005; Norman, 2005; Wallace, 2000). This form of control that does not allow a human to interact directly with the machine any more leaves room for the application of shared control concepts to a domain that, although cites shared control, is really based on supervised control with little

to no shared interaction (Norman, 2005).

Fuzzy Mediation is an innovative framework that allows two controllers to interact in real-time while guiding an object, with a dynamic allocation of the weight of control. The majority of information fusion algorithms involve a static method of evaluation set at the time of system design (Kokar et al., 2004). The concept introduced in this work is based on a dynamic concept of mediation, assuming that the two controllers can interact and learn, thus the need for a reward system shifting control to the learner, as its knowledge allows for actions that resemble the trainer's inputs.

As it would be impossible for a vehicle to be controlled simultaneously by two different controllers, we looked at concepts of information fusion for a solution. The fuzzy controller performs three distinct operations. The first is the analysis of the inputs to determine the closeness of control; it then performs a revision of the weight of control between the expert and the novice controller; finally it computes the value of the single output.

3.1.1. Analysis of the Inputs

Cantorian set theory leaves little room for gray areas; a room's temperature can be classified as hot, medium or cold. This system does not take into consideration the possibility of the same room being perceived as comfortable or chilly by two different persons. Fuzzy sets provide a solution that takes into consideration values that fall within multiple sets (Zadeh, 1965). These kinds of sets can be utilized within the field of information fusion applied to a situation of multiple controllers trying to interact with a vehicle by a means of comparison.

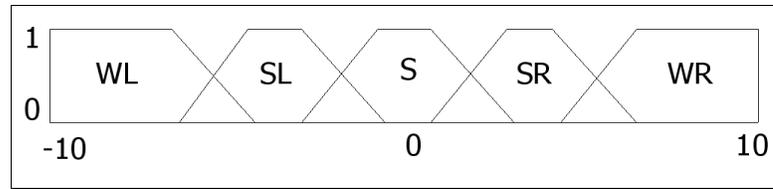


Figure 14: Fuzzy sets used for classifying the deviation between inputs

The analysis of the inputs coming from the two controllers aims at understanding the distance between the values. The input coming from the expert user is mapped to the center of the range $[-10, 10]$, as shown in Figure 14.

The range of $[-10, 10]$ represents the highest possible deviation between the input of the expert and the one of the novice. After the input of the expert becomes the center of this domain, we calculate the difference between the value of the original input of the expert and the one of the novice. This value is then also mapped to the domain shown above.

When the distance between the inputs is mapped, it will fall within one or two sets that span over the range, as shown in Figure 14. The five sets we deal with are: WL (Wide deviation to the left), SL (Slight deviation to the left), S (Similar), SR (Slight deviation to the right), and WR (Wide deviation to the right). In fuzzy set theory, the value can belong to a set with a certain degree of belonging. Such degree of belonging is calculated by a set of simple membership functions. In our case, we use a simple membership function, also shown in Figure 14.

The application of a linguistic modifier to the deviation between the controller inputs also keeps in consideration the degree of belonging of the difference to the different sets. It is important to note that a value may fall completely (degree of belonging = 1) within one set, or it may belong mostly (0.80) to the set of deviations deemed as Similar (S), and partly (0.20) to the set of Slight deviations to the left (SL).

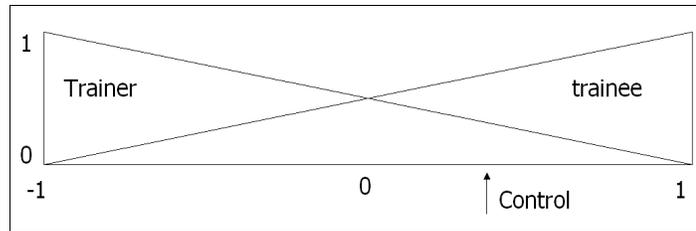


Figure 15: Fuzzy sets that regulate the balance of control

3.1.2. Revision of the Weight of Control

Fuzzy Mediation sees the fusion of the inputs of the expert and the novice as a balance between the two. The more the novice performs similarly to the expert, the more control will shift in favor of the second controller. Likewise, the more the control of the novice differs from the one of the expert, the more control will shift back towards the expert.

Given this preamble, this second part of the Fuzzy Mediation algorithm analyzes the linguistic modifiers applied to the deviation of the inputs during the first phase. Control is mapped to the range of $[-1, 1]$, where a control weight of -1 identifies a control fully in the hands of the expert, a value of 1 instead refers to control managed by the novice. Figure 15 shows the visualization of this concept. As we apply the concept of fuzzy sets to this section of the algorithm, a value in the range $(-1, 1)$ identifies a control that is mixed in a certain proportion between the expert and the novice. The arrow in Figure 15 shows a possible weight of the mediation of control between the trainer and the trainee. At the beginning of the simulation the weight has a value of -1 .

The shifting of the weight from one controller to the next occurs in a linear fashion, with increments or decrements of 0.2 points on the range $[-1, 1]$ presented earlier based on Mamdani inference rules (1975) as a generalization of the if-then inference in fuzzy

logic terms (Tran, Jain & Abraham, 2002). When the classification of the distance between inputs is analyzed, there are several actions that can be taken. If the inputs are classified as similar, control is given more to the novice (*Weight of Control*+0.2). If the deviation between inputs is slight, then control stays unvaried (*Weight of Control*+0). If the deviation is wide, more control is given to the expert (*Weight of Control*−0.2). In the case of a distance between inputs that belongs to two sets, then we will multiply the degree of belonging to each of the sets to the action associated with that particular set.

Using the example given earlier, a value that belongs to the set S with membership 0.8 will receive an increase in control of 0.2 (the standard increment) multiplied by the membership value, which means an increase of the weight of control of 0.16. The same value also belongs to the set SL with membership 0.2. The action associated with a deviation that is classified as slight is a movement of 0 of the balance between controllers, so the action for this set is calculated by multiplying 0.2, the membership value, to 0. The addition of these values, 0.16 and 0, shows the overall shift in control, which is of 0.16 in favor of the trainee.

3.1.3. Calculation of the Single Output

After the weight of control is updated, we need to calculate a value that will serve as a single input stemming from the original inputs of the two controllers. For this computation we need to refer to the original values. Equation 1 regulates the third section of this algorithm.

$$MO = \mu_T \times E_I + \mu_I \times N_I \quad \text{Equation 1}$$

where MO symbolizes the mediated output, μ_T refers to the membership value of

the weight of control to the Trainer set (Expert) and μ_t refers to the membership value of the weight of control to the trainee set (Novice). The inputs of the two controllers are represented by E_I for the Expert's input and N_I for the Novice's.

In the case of a driving simulator, we may have an expert applying a turn of 15 degrees to the right and a novice applying a turn of 25 degrees. If the weight of control has a value of -1 , the mediated output will have a value of 15 degrees to the right. Likewise, if the weight of control has a weight of 1, the mediated output will be of 25 degrees to the right. If the weight is anywhere in between, for example $\mu_T=0.5$ and $\mu_t=0.5$, the mediated output will be of a 20 degree turn to the right.

3.2. Demonstration of Concept

This section introduces the concept of Fuzzy Mediation applied to a simple simulation. For these experiments we use two sometimes very different equations to analyze the behavior of the overall output. Each environment will explore the behavior of two functions in different ranges of the X value, in order to explore the region where the two behaviors first seem to converge, then touch, then depart from each other once again.

When we put Fuzzy Mediation to work we expect to see the computed output to resemble the behavior of the expert controller when the novice is really far off. As the behavior of the novice starts resembling the one of the expert, then the mediated output will converge towards the one of the novice. As the two controllers perform in similar manners, the computed output will still resemble the one of the novice. As the novice starts performing differently from the expert, then the computed output will start converging again towards the expert. Finally, the expert will retain control as the novice

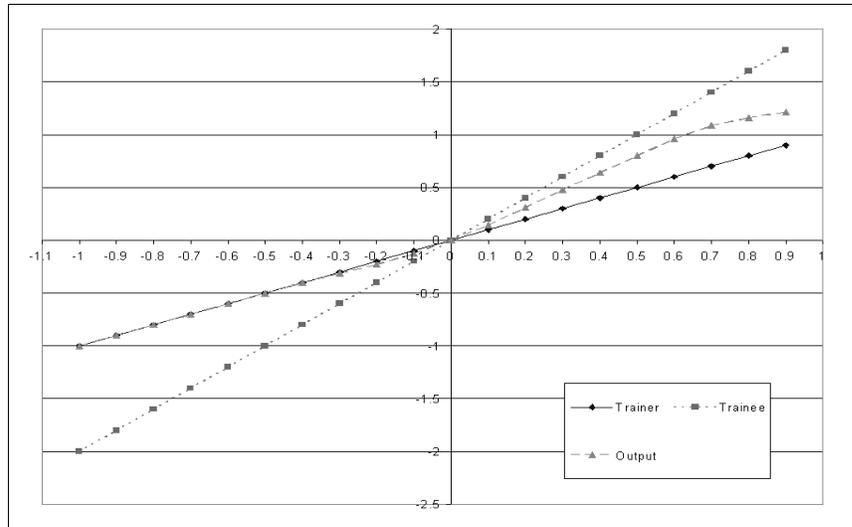


Figure 16: Simulation with two linear functions

performs quite differently.

A successful demonstration of the concept would show the computed output to resemble closely the expert controller's behavior first, then control will shift either partially or completely to mimic the novice's, to then shift again and become similar to the expert's once again. It is important to mention that no controller tries to learn from the other, neither the expert nor the novice are programmed to adapt to the situations. The two controllers will simply compute the output Y for the given value X in according to the preset function that powers that particular controller.

The first simulation that we performed focuses on two linear paths that intersect at some point, without any deviation on the side of the novice to try and resemble the actions performed by the expert, as shown in Figure 16. As we can see, the lines that represent the expert and the novice intersect at one point only, and they do not converge past the intersection. The output that is produced by the Fuzzy Mediation engine overlaps the path taken by the expert, since the path traced by the novice is quite distant. As the

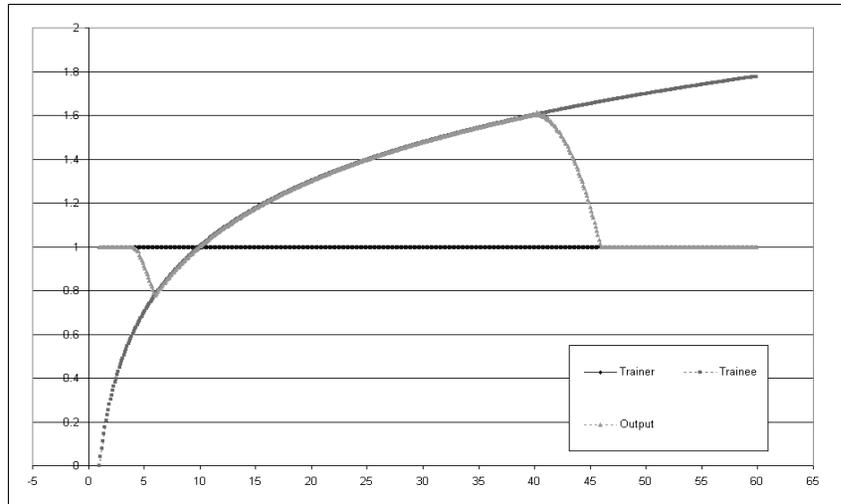


Figure 17: Simulation with an expert linear function and a novice logarithmic one

controllers perform actions that are closer, the mediated output starts shifting towards the novice's line. After the intersection, the output of the fuzzy controller shows that control is shared somewhat equally between the expert and the novice for a few steps, but then it leans towards the expert's output once again. This is due to the fact that the novice's controls deviate widely from the ones of the expert, thus taking control away from the first and giving it more to the second. The output created by the fuzzy mediator does not adhere much to the line of control carried out by the novice because the time during which the output of the novice and the ones of the expert were somewhat similar is rather short. In cases where the lines would not intersect, the control will never move away from the expert's line.

In the second simulation, shown in Figure 17, we have taken into consideration a situation where the expert is controlling the vehicle in a linear fashion, and the novice performs a motion trying to align to the direction of the expert in a logarithmic manner. Also in this case, when the simulation starts, the output of the fuzzy controller is

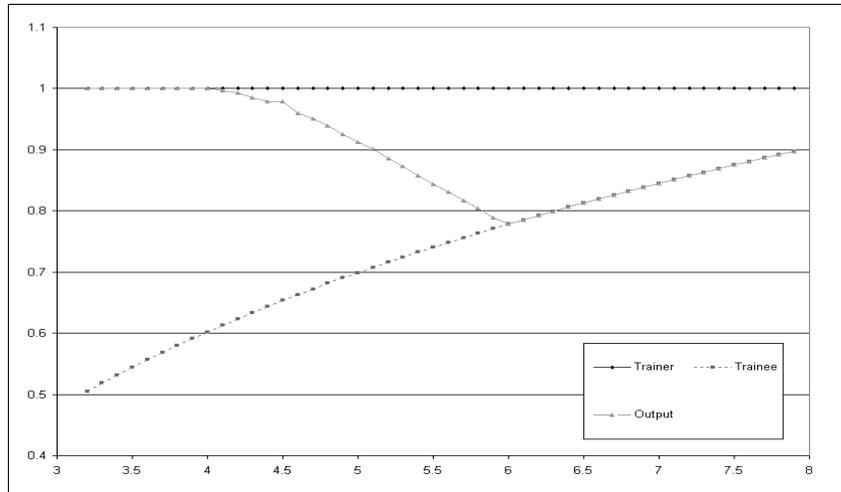


Figure 18: Detail of the area in Figure 17 where control shifts to the novice

overlapping the control of the expert. As the novice's output gets closer to the expert's, the output of the Fuzzy Mediation shifts. This time, the two values are close enough for quite a few cycles of comparison within the Fuzzy Mediation system, where each time the two values are considered similar, giving more control to the novice.

Figure 18 shows a closer look at the area of the plot where the output of the Fuzzy Mediation engine starts moving towards the novice's controls. We should note that the difference between values, at that point, is only of 0.4 units. At this point, control is in the hands of the novice completely.

After the lines intersect and then diverge, control will be retained by the novice, at least until the fuzzy mediator perceives a difference between the two inputs that is significant enough to start shifting control back in favor of the expert. This happens when the two values diverge by more than 0.6 in this simulation. The area where control starts shifting back to the expert is highlighted in Figure 19.

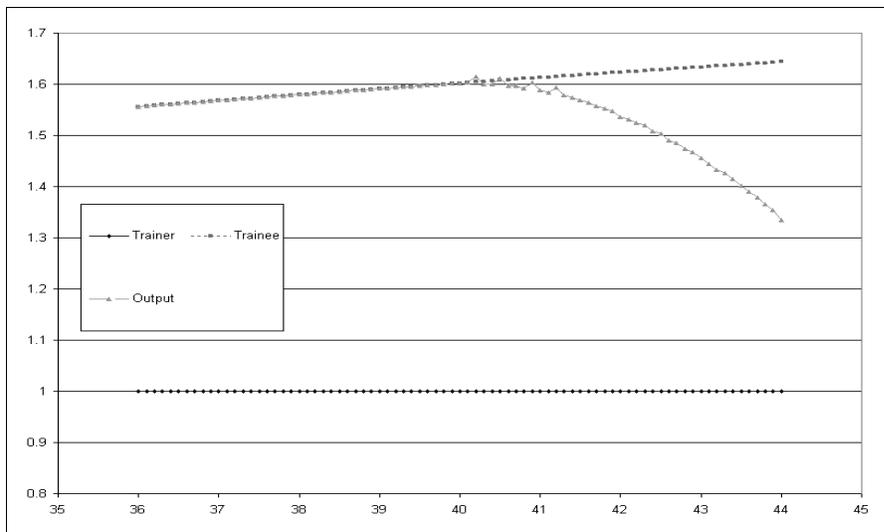


Figure 19: Detail of the area in Figure 17 where control shifts to the expert

In Figure 20 we can see yet another simulation. In this case, the expert and the novice perform actions that are at times similar, and then diverge to become similar again later in the simulation.

In this particular situation, the actions are periodic, so that we can see that the Fuzzy Mediation engine performs consistently. The expert's control is represented by wave closer to the X-axis. The control is initially given to the expert and is slowly shifting towards the novice. As the novice's performance deviates from the expert's, the expert's

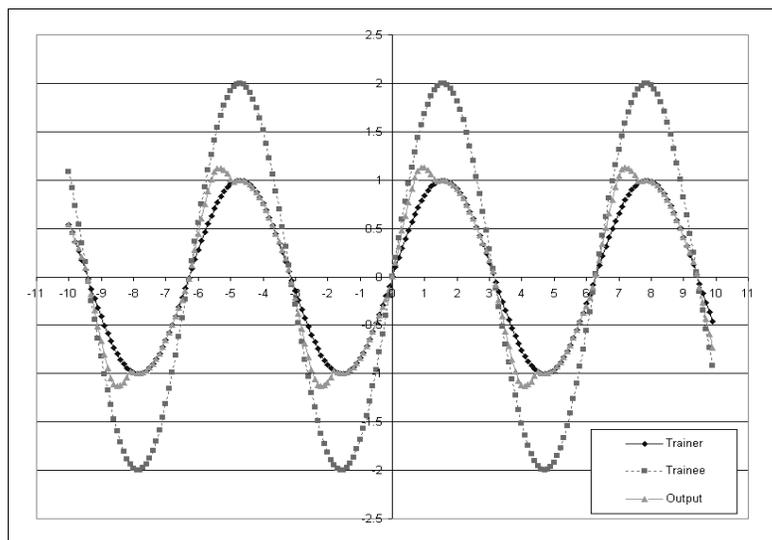


Figure 20: Simulation with two sine waves

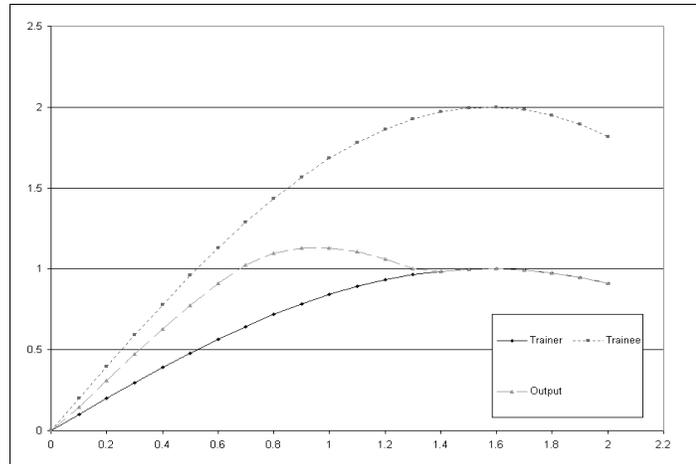


Figure 21: Detail of one of the areas in Figure 20 where control shifts to the expert

input starts weighing more and more in the overall output of the fuzzy controller. Figure 21 focuses in on the area of Figure 20 where control shifts back towards the expert.

3.3. Summary

In this chapter we introduced the concept of Fuzzy Mediation. First, we looked at a very high-level perspective of the niche in which this dynamic algorithm will fit into the greater scheme of things. Then we looked at the more technical aspects of the algorithm, with a detailed breakdown of the three main steps to the fusion of the inputs. Finally, we reviewed a set of very simple experiments that served as a demonstration of the concept, illustrating how the overall behavior of the mediation algorithm affects the final output.

In the next chapter we will review several concepts that will define the framework used for the more elaborate experiments used to show the effectiveness as well as the importance of Fuzzy Mediation as an innovative and dynamic extension to the field of Information Fusion.

4. Prerequisites

In this chapter we will review some elements that will be used during the tests performed to evaluate the performance of Fuzzy Mediation. First, we will take a look at the two simple simulators that were created to analyze the performance of artificial and/or human agents interacting. Then, we will review the performance of agents based on the Wang-Mendel algorithm. In this section we will also review the performance of these agents as they face different training routines in order to evaluate the overall appropriateness of the use of this controller in these simulations.

It is important to note that a thorough evaluation of different training methods on the Wang-Mendel algorithm was never performed before, thus the material shown in that section is also the result of our studies. Others did not study this particular algorithm because of its poor performance when the number of dimensions to control grew. It does adapt quite well to 1-Dimensional domains of control, such as the one used in the following experiments.

4.1. Simulation Environments

Since Fuzzy Mediation was still an uncharted territory, we decided to study the behavior of this algorithm by using custom-built simulators. We required environments that would accommodate several kinds of simulations, thus they needed to be extremely flexible. The two simulators that were built for these experiments are described in great detail in these next two sections. Each environment was created to test for different metrics and concepts. As the environments were also new, we conducted some preliminary tests, especially for the line follower, to document the behavior of agents

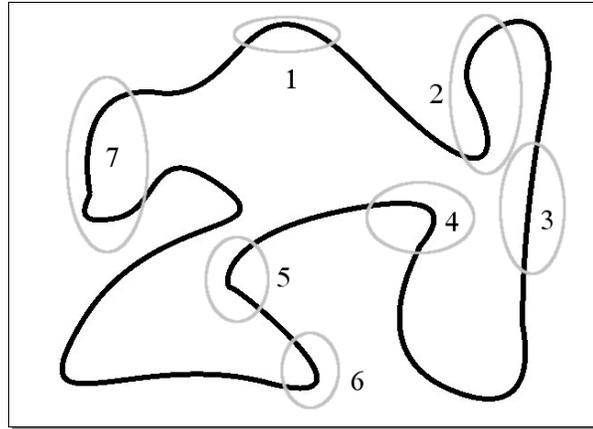


Figure 22: Path of the simulation

based on artificial intelligence when placed in situations like the ones we created.

4.1.1 Simple Two-Controller Line Follower

The environment used for these experiments is a simple agent that follows a line. Figure 22 shows the pattern that was used. The pattern contains only the white background and the black line. The areas that have been highlighted show the sections of most interest. Section 1 was created to see the behavior of the agent in a mild turn to the right; section 2 instead simulates a sharp turn to the left followed by a moderate turn to the right. Section 3 mimics a straight path. Section 4 reveals a tight turn to the left, just like section 5. Section 6 is another tight turn to the right after a straightway and finally section 7 shows quick changes in direction.

In these experiments we use different levels of control. When we want to create an environment when the two inputs need to be closer in order to shift the weight from the expert to the novice we simply need to set the boundaries of the sets shown in Figure 14 to tighter limits. The three levels of control we use are the following: Tight control,

	<i>Tight</i>	<i>Moderate</i>	<i>Loose</i>
WL	$(-\infty, -\infty, -5, -3)$	$(-\infty, -\infty, -6, -4)$	$(-\infty, -\infty, -8, -6)$
SL	$(-5, -3, -2, -1)$	$(-6, -4, -2, -1)$	$(-8, -6, -4, -2)$
S	$(-2, -1, 1, 2)$	$(-2, -1, 1, 2)$	$(-4, -2, 2, 4)$
SR	$(1, 2, 3, 5)$	$(1, 2, 4, 6)$	$(2, 4, 6, 8)$
WR	$(3, 5, +\infty, +\infty)$	$(4, 6, +\infty, +\infty)$	$(6, 8, +\infty, +\infty)$

Table 1: Limits of fuzzy sets used for classifying the deviation between inputs

Moderate control and Loose control. Table 1 shows the values associated with each level of control. Each set carries four values, (OL, IL, IR, OR), where OL refers to the outer left boundary, IL to the inner left, IR to the inner right and OR to the outer right.

The simulated agent is composed of a central unit that contains sensors. The sensors check the terrain in front of the agent for color. The sensors can either pick up white, which is the background, or black, which is the line. The sensors are arranged on a probe that scans the range $[-45, 45]$ in front of the agent at 5-degree intervals. Figure 23 shows an image of the agent that is following a line. The light gray is the body of the agent while the dark gray spots in front of it represent the range of action of the sensors.

The sensors communicate to the agent the color of the terrain at each angle. Then the agent will group together the angles that recorded a reading of a line and calculate the average. Such average will be analyzed by the agent, which will select the new heading.

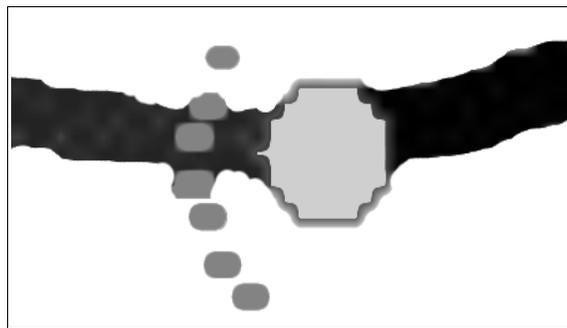


Figure 23: Diagram of the simulated agent with its sensors

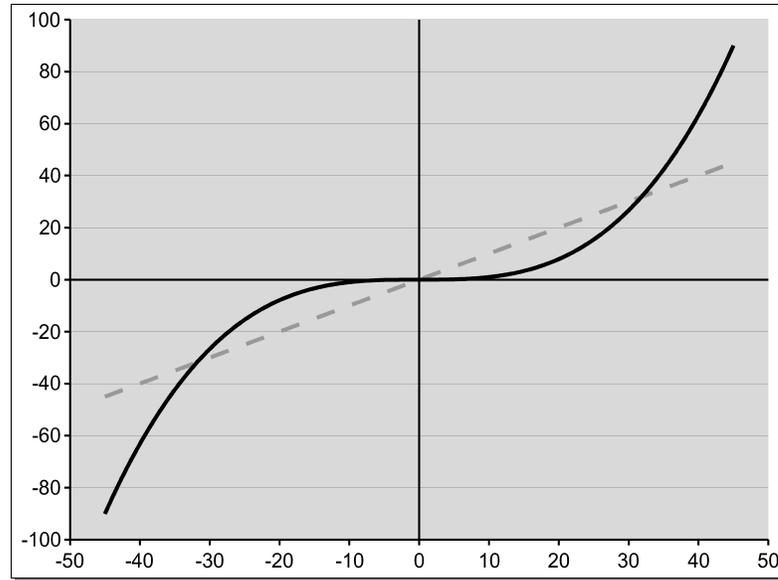


Figure 24: Control functions, expert (dotted) and novice (solid) controllers

An agent can perceive changes in direction up to ± 45 degrees.

In order to simulate the behavior of agents we assigned them a preset behavior that allows them to navigate successfully through the pattern selected. In order to simulate an expert agent and a novice one, we chose equations that are slightly different. The typical expert is represented by a simple linear function. When the agent receives the reading from the sensors, the difference in heading is applied directly to the heading, so if the sensors read that, in order to follow the line the agent needs to apply a 15-degree turn to the right, the agent will perform a 15-degree turn to the right.

The simulation that represents the case of a novice is powered by an agent that relies on the cube of the difference normalized to the range $[-90, 90]$. The two equations that are used to drive the agents are reported in Figure 24. The dotted line represents the behavior of the expert and the solid one the behavior of the novice.

This interpretation of the agents shows an expert that acts as expected, with a linear response to the situation. The novice instead reacts more slowly only to overcompensate

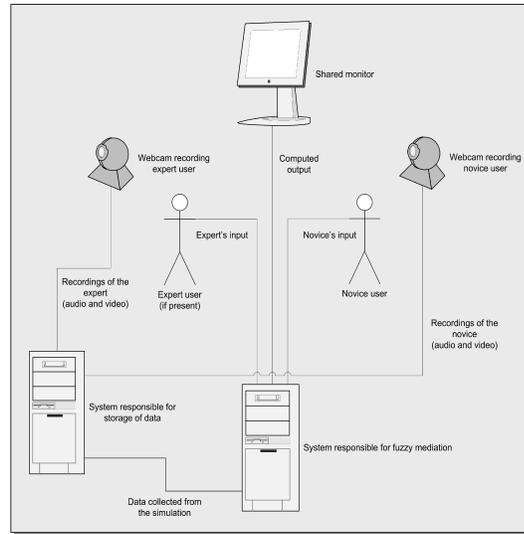


Figure 25: Experimental setup for the simple flight simulator

as the deviation required in order to remain on track increases. We also carry out other simulations where the difference between the inputs of the expert and the one of the novice are very different.

4.1.2. Simple Flight Simulator

This experiment aims to study the change in the rate of learning as the novice (participant) will interact with the simulation first-hand, while receiving guidance from an expert user.

The setup of the system the participants interacted with is described in Figure 25. The participant used one or both the joysticks to maneuver the simulated aircraft through the task. A list of the hardware used is reported later in this document.

Each participant was subjected to the following four activities:

- Administration of the pre-study questionnaire
- Simulations

<i>Step</i>	<i>Task</i>
1	Attempt 1
2	5-minute break
3	Attempt 2
4	5-minute break
5	Attempt 3

Table 2: Script for simulations studying shared control

<i>Step</i>	<i>Task</i>
1	Attempt 1
2	5-minute break
3	Attempt 2
4	5-minute break
5	Attempt 3
6	Break
7	Attempt 4

Table 3: Script for simulations studying bi-manual control

- Administration of the post-study questionnaire
- Debriefing

We performed two sets of experiments based on this script. The only variation between the two will be within the “Simulations” section. The scripts followed by the two tests are reported in Tables 2 and 3.

Each attempt within the “simulations” section involved the participant to interact with the simulation. Figure 26 reports a diagram showing the typical environment the participants interacted with. They were required to maneuver the airplane from the

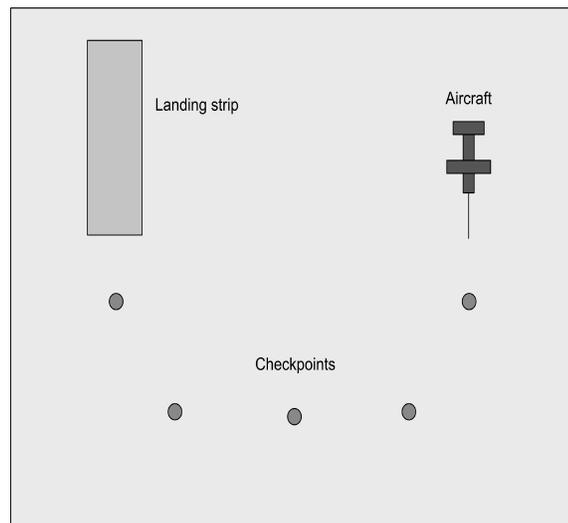


Figure 26: Diagram showing the simulation environment

starting position shown in the figure to the runway. Along the way, the participants were lead by checkpoints, reported in red. The participants were required to pass over the checkpoints in a particular order, starting from the one closest to the aircraft, then the one following and so on. Upon passing the checkpoint, the red color turned into green, signaling to the participant that the goal of passing over that checkpoint was achieved. The final goal of the simulation was reached as the aircraft got to the landing strip from the bottom, as it should happen, if the participants follow the order of the checkpoints correctly. If the participants did not touch checkpoints, they were asked to proceed to the next one, without worrying about the ones missed. It is also important to note that the aircraft was subjected to an external force, acting as an aircraft would in a moving fluid, such as a flying on a windy day.

The interface for the weight of control was represented by a square above the airplane. Such square would assume different colors as the mediation was dynamically changing. The more the color of the square was white, the more the weight was leaning towards the expert controller. In the case that the weight of the mediation was leaning towards the novice, the color would become brighter green. No other indications were given to the users.

As stated earlier, we performed two kinds of experiments, with slightly different goals. The main target of the simulator was identical, the analysis of the performance of users. The first round of tests was performed to study how joint control of the aircraft affects the performance of individuals in this setting. The second batch of tests instead focused on the performance of a single individual and how this particular individual was

able to transfer motor knowledge from their dominant hand to their non-dominant one using Fuzzy Mediation techniques. We will now describe both tests in detail.

In the first simulation, the first attempt required that the participant would navigate through the simulation without any knowledge of how to control the aircraft or if any external forces were applied to it. This was regarded as an exploratory run where the participants were exposed to the concept, without any formal explanation.

The second attempt varied between the control group and the experimental one. The participants belonging to the control group were required to go through the simulation once again, without any instructions. This is to simulate a typical user who prefers to learn a motor skill by exposure to the task, instead of any type of learning from another source, such as an expert or an instructions manual. The experimental group instead was required to perform this attempt with the help of an expert user. The expert user mentioned above is a person who is familiar with the simulation and can successfully perform the task. Fuzzy Mediation played a vital role in the fusion of the inputs for the creation of a single output that would regulate the flight of the aircraft. The expert user and the novice will be required not to talk about the principles that govern the navigation of a vehicle in a fluid, such as an airplane flying in windy conditions.

The third attempt wants to record the performance of the two groups after the administration of the experimental variable. We expected to record a better performance for the experimental group, as the participants should now be familiar with the

environment of the simulation.

The second experiment instead was structured in a different way. During the first attempt, the user tried to navigate through the simulation using the dominant hand. This ensured to give us a good baseline to compare the following performances. The second attempt then established the baseline for the performance of the non-dominant hand.

For this set of simulations, the third attempt was the one that constituted the differentiation between the experimental and the control group. As stated earlier, in the second set of tests we are trying to analyze how well Fuzzy Mediation accommodates to the transfer of motor learning from the dominant hand to the non-dominant one. In this run people who belonged to the control group were asked to simply perform the simulation with their left hand, as a training run. In the experimental group, the subject was given two controllers; the one connected to the expert's side of inputs was to be controlled by their dominant hand and the one that functioned as novice was controlled by the non-dominant hand.

During the last attempt the subject was asked to perform the task required with the non-dominant hand to check if there was any type of improvement based on the different kinds of training.

It is important to have breaks between the attempts as to ensure that the performance of the participants is not linked to a simple habituation of the muscles to the task, but it involves long-term retention (Kandel, Schwartz & Jessell, 2000).

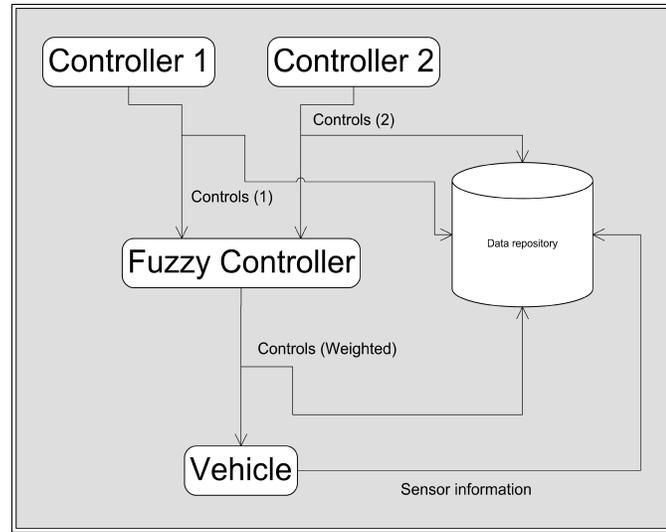


Figure 27: Information flow architecture for the simulations

The questionnaires that were given to the subjects are reported in the Appendix. The debriefing included a simple explanation of the motion of a body in a moving fluid.

The setup of the experiment allows the recording of the following elements:

- Action of the participant on the joystick
- Action of the expert user on the joystick (if present)
- Position of the aircraft
- Reactions (video camera)

Figure 27 outlines the architecture used, and at what point within the simulation the data was recorded. Controls (1) refers to the interaction with the joystick of the expert user, should there be one. Controls (2) instead refers to the interaction with the joystick of the participant, the novice user. Controls (Weighted) refers to the value that will point the aircraft in one direction or another. Sensor information, finally, refers to the position of the aircraft. As the users will be able to view the aircraft's position as well as the distance

and the deviation from the checkpoints and the desired heading, we need to record that information for the purposes of analyzing the simulator data at a later time.

The hardware we use is listed here:

- Logitech QuickCam STX (2)
- Logitech Attack 3 Joystick (2)
- Standard PC for the simulation
- Standard PC for storing the data gathered as well as the video and audio recordings
- 17" CRT monitor that will display the simulation

It is important to mention that the joystick was selected because of its architecture.

The Logitech Attack 3 Joystick is one of the few joysticks that are designed for accommodate right-handed as well as left-handed users.

The user interacted exclusively with the joystick, as s/he will simply play the simulation. There was no interaction with any other software on the computers to be used.

4.2. Learning in Agents Based on the Wang-Mendel Algorithm

Among many algorithms that emulate human intelligence, the Wang-Mendel algorithm (Wang and Mendel, 1992) is a very efficient one when applied to problems of vehicular control. Its first application, reported in Wang and Mendel (1992), employs this method to a truck “backer-upper” problem, where the target of the simulation was to

place a truck into a virtual environment, and let a controller back it up to its loading station. In their experiments, Wang and Mendel used an expert to teach the program how to “back up” into the loading dock, and then let the application perform. They found that their algorithm was extremely efficient in solving this problem.

The Fuzzy Associative Memory, or FAM, stores the relationship that exists between one or more events and a consequence. The details of this process are reported in (Wang and Mendel, 1992). One important aspect of this algorithm is how the system reacts when multiple events are associated with different consequences within the training pairs.

When multiple examples lead to similar associations, or should the example show the FAM that an association is different from what has been previously stored, the algorithm takes into consideration the strength of the association already stored. If the new association should have a strength that is higher than the one already present, then the association will be changed, if necessary, and the new strength of the rule will be stored. Should the strength of the training example be lower than the association already stored, then the example will not be considered. The importance of an association is calculated by multiplying the membership values of the fuzzy sets that have the highest ownership of the input (event) and output (consequence) for the training pair.

For this simulation we will use this algorithm given its ideal use shown in the original work. We will simulate the control of an agent. Such control will be represented by the equations shown below.

$$Y = X \quad \text{Equation WM1}$$

$$Y = |X| \quad \text{Equation WM2}$$

$$Y = \sin(X) \quad \text{Equation WM3}$$

$$Y = X^2 \quad \text{Equation WM4}$$

These four equations will create training examples for the FAM to learn. In order to train the memory, we will use several training algorithms, discussed in further detail shortly. After each training session, we will test the memory by asking it the Y for a given X. We will proceed in an ordered manner, with X starting at the lower boundary of the range and increasing until it reaches the higher boundary.

4.2.1. Training Algorithms

In order to carry out a simulation that is realistic, we have analyzed several training scenarios that would mimic an autonomous agent with imprecise knowledge at first that would improve as the operations continue. The training algorithms are quite different, and they are: “Greedy”, “Random”, “Core” and “Left to Right”.

Greedy training. This type of training will feed to the fuzzy system random numbers until the FAM is just filled. The accuracy of this type of learning is highly dependent on the set of training numbers. If the set happens to include values that are right around the cores of the sets, then this algorithm will perform very well, otherwise the rules may not be the most suitable.

Random training. This training style is the one where the memory is presented with a fixed number of training pairs. If a rule has not been found for a particular association of sets, then the system will output the default value stored in the FAM. There is no assurance that the best example will be used for a certain association.

Core training. This type of training feeds to the system the values of the cores. Theoretically, once a rule reaches a point where the membership of the input value to a set is 1 (the value belongs only to one set), then the rule will not be replaced by any other one, because that sample is the best available for that set. This type of training is significant because we will be able to compare how many training sets will be necessary to achieve results similar if not equal to the results that we can achieve if the training set was the set of all core values for the input and output sets.

Left to Right training. This final algorithm for training the memory is used as control-type training. The training sets are given to the memory in an ordered manner, starting from the lower boundary of the range of the training set and ending at the high boundary. This type of training ensures that the FAM is given training pairs that cover the entire range of the simulated environment.

4.2.2. Analysis of the Various Training Algorithms

In order to evaluate the performance of the Wang-Mendel algorithm, we analyzed the performance of a hypothetical agent trained using the algorithms described above. Also, we wanted to evaluate how accurate the performance of an agent based on this algorithm is as we increase the number of sets within the decomposition, or granularity, of the input and output sets associated with the fuzzification and defuzzification operations. We chose the average error as index in order to evaluate the performance of the various agents (Vincenti & Trajkovski, 2006).

After the initial training that the agent receives in the form of (X, Y) , where X is the input and Y is the expected output – computed by equations WM1 through WM4,

<i>Eq.</i>	<i>Left-to-Right</i>	<i>Core</i>	<i>Greedy</i>	<i>Random</i>
WM1	7.61359×10^{-15}	7.61359×10^{-15}	7.61359×10^{-15}	7.61359×10^{-15}
WM2	9.91578×10^{-15}	9.91578×10^{-15}	0.126954513	9.91578×10^{-15}
WM3	0.194910328	0.189261003	0.287849826	0.203842502
WM4	0.539344649	0.458899104	0.747058728	0.561321949

Table 4: Average error for Equations WM1 to WM4, versus training algorithms

depending on the scenario – the agent is then asked to return a Y value for every X it is presented. The calculations are made on the absolute difference between Y_{EXP} , the Y calculated by evaluating X by means of the equation used for training, and Y_{FAM} , the Y calculated in according to the Wang-Mendel algorithm and based on the trained agent. Each agent is tested using 201 equally spaced data points. Table 4 shows the average error for all the simulations performed. The values show the average error for all agents set with a granularity between 7 input/output sets and 21 input/output sets in odd increments.

The data show that the agents learn effectively in situations where the relation between an X and a Y is linear, with no changes in the trend of the relationship, as in equation 1. The agents show a slightly greater error in a system where the linearity between X and Y is still present, but there is change, as the one associated with equation 2. We see the greatest values for error in the tests performed using equations 3 and 4. Any type of training will produce a greater error than in the case where the relationship between the two variables is linear.

Among the training algorithms we can see that the “Greedy” training is the method that yields the greatest error, because the agent starts working as soon as there is one association for each input/output region. This allows for some rules that have a very weak importance to stay in the system, and not be overwritten by stronger rules. The number of

examples shown to the agent is highly dependent on how quickly the random examples cover the entire span of the input/output range.

“Random” training yields results with lower error. It is important to note that this is due to the fact that some weak rules are replaced by stronger ones found during successive observations. In this case the agent is shown 250 training samples, even if all the input/output rules have already been filled.

The “Core” training shows the smallest error for the entire set of simulations. Theoretically, the strongest associations are the ones where the membership of a value to a set is closest to 1, or full membership to a set. In this particular system, the core of a TPE contains only one value, and that is the value that was used to train the agent. There is no other example that will replace the given example, because the example given already maps to a membership of 1. In this case the agent is shown as many examples as the number of fuzzy sets that define the input/output range, thus the agents were shown between 7 and 21 examples, in odd increments.

The “Left-to-Right” training is considered a control type training because we train the agent in a methodical manner with examples starting from the lower boundary of the range and reaching the higher boundary. We have found that the error for this type of training is slightly higher than the “Core” training. The explanation lays in the fact that the example shown to the agent may not have the highest degree of importance, thus is not the best association.

Finally, the simulations related to the last situation, the one governed by equation WM4, only reinforce the findings associated with what was discussed for equation WM3.

The next aspect to analyze is the relationship between the average error, the type of

	<i>Left-to-Right</i>	<i>Core</i>	<i>Greedy</i>	<i>Random</i>
7	0,627506065	0,675936524	0,650292194	0,650422003
9	0,237706944	0,237706944	0,400731338	0,230864971
11	0,197668687	0,171135865	0,252355939	0,241180776
13	0,139887356	0,138705757	0,246697204	0,154017493
15	0,108492532	0,115646522	0,249215771	0,107077011
17	0,074813892	0,064634848	0,141544446	0,096659755
19	0,112382012	0,061053242	0,167742751	0,075847614
21	0,060825138	0,04926832	0,194218968	0,074670394

Table 5: Average error of Equation WM3 given different granularity and training

training, and the number of sets associated with the input/output domains of the agents.

Table 5 shows the results relative to this aspect of the experiment.

The results of this analysis are consistent with what would be expected. The higher the granularity of the input/output domains, the more accurate the agent's performance is. For this simulation we analyzed the experiments that were conducted using equation WM3. We notice that in most cases, no matter what the training algorithm was, the average error is reduced as the number of fuzzy sets is increased. This is not true only for the "Greedy" training, as it reaches 21 fuzzy sets. The higher error can probably be traced back to the fact that the agent was shown only one example for some of the associations, without the possibility of replacing such bad example with another one with a higher significance. It is important to note that an agent performing with 21 fuzzy sets and that has been trained with a "Core" algorithm has the lowest error rate of all the simulations reported above.

In addition to the tests performed to show how the Wang-Mendel approach operates, we also explored the dynamics of this approach to an artificial agent when faced with yet another equation, shown below:

$$Y = X^3$$

Equation WM5

	<i>Left-to-Right</i>	<i>Greedy</i>	<i>Random</i>	<i>Core</i>
3	248.7313433	251.2437811	248.7313433	248.7313433
5	44.86827861	44.86827861	44.86827861	44.86827861
7	35.96318408	35.96318408	35.96318408	35.96318408
9	64.41693532	56.87746269	67.75329353	54.64260697
11	34.91802985	57.92933333	42.84458706	34.91802985
13	23.69535655	38.35538972	37.1592869	23.69535655
15	21.78149538	40.08151812	27.26113149	21.78149538
17	17.52935323	27.6420796	24.89635821	17.52935323
19	20.6328911	25.12282477	23.98280818	20.6328911
21	14.67482587	26.8941592	19.45707463	14.67482587

Table 6: Average error for Equation WM5, various granularities and training

The agent was exposed to the same training routines described earlier in this work. Table 6 shows the average error reported in the experiments given a different granularity and different trainings. The random training for this set of experiments was performed with 250 items.

It is interesting to note that, with this equation, the FAM performs differently from previous tests. As the number of sets increase from 3 to 7, the performance of the agent improves. When the number of sets is increased to 9, the error increases noticeably for all types of training. The performance of the agent will then improve until it reaches the best performance with a granularity of either 19 or 21 sets. The reason for which, during *Greedy* training, the agent performed better with a granularity of 19 sets as opposed to 21 sets is probably due to the fact that the core of the decomposition sets fell on points that would represent more accurately the equation to be learned. This is due to chance rather than a planned study of this particular granularity.

It is important to also note that the agents tested with Equation WM5 included FAMs with granularities as low as 3 sets as opposed to previous experiments. This is

because the focus was shifted from the behavior of an agent given a certain training routine to the actual performance given extreme situations. Agents with a very low granularity will be used in later experiments to evaluate the performance of Fuzzy Mediation in situations where the agent's best performance still does not allow it to perform the desired task within our simulation.

At last, we can see that the average error for different training algorithms reflect the findings discussed earlier in this section. This is especially true as the number of sets that the agent uses increases.

4.3. Summary

In this chapter we reviewed several components that will be vital in understanding the operations of agents as well as humans in the experiments described in the following chapter. First, we reviewed the simulation environments built especially for this algorithm. The first will be used in tests involving artificial controllers, both static and A.I.-based. Then we introduced the simple flight simulator that human controllers will use to carry out a section of the tests. Then we looked into much detail of the operations of agents based on concepts based on the research of Wang and Mendel (1992). Since this type of A.I.-based controller was never analyzed carefully when exposed to different types of learning, we conducted some experiments to establish a baseline that we could compare to the tests reported in the next chapter.

5. Experiments and Results

As we saw earlier, Wasson and Gunderson (2001) state that it is important for communications to occur between controllers in a collaborative control situation, and that the situation only gets more sticky when humans are involved in the loop. The algorithm proposed in this work creates a framework that is universal to systems where both agents are automated, when one automated agent interacts with one human agent, or when two human agents are involved. This eliminates much overhead in the creation of a communications system for multiple agents to understand each other.

Generally, the application of a controller to an agent involved two phases: Training and Performance. The controller can direct the agent only after training has occurred. With Fuzzy Mediation, there is a constant dynamic shift between the training and performance phases, making it possible to let the controller direct the agent from the very beginning. As the permanence in the Training stage fades, Fuzzy Mediation becomes a metric that analyzes how well the trained agent is performing.

An agent's performance may improve or get worse as it learns from the expert. This is due to the fact that some examples may yield associations with better performance but lower weight that then will be replaced with associations with worse performance but higher weight. In the agents we use, the weight of a rule is dictated mainly by how close to the core of the classification set a value is. Other artificially intelligent controllers may react differently to this framework. The limitations of the controllers are mainly due to the performance of the Wang-Mendel algorithm, which behave differently given different

granularity Fuzzy Associative Memories.

The novice is continuously learning from the trainer, even if the simulation is in the hands of the trainee. This simulates an environment where the trainee is getting feedback from the visual display that shows that the actions the trainee performs are similar to the ones of the trainer, thus hopefully reinforcing the associations and perhaps producing rules with higher weight.

The only time when the novice is not observing the expert is ideally when “emergency” mode kicks in. This mode allows the expert to override Fuzzy Mediation and gain full control. This is a safety mechanism that allows the expert to maneuver the agent in situations where it is crucial that accurate control is exercised without the influence of any similar controls the novice may show by accident. This particular mode is not shown in these simulations, as we were focusing only on the interaction component, and we did not have any robotic agent to salvage in the case of crashes of the system.

Also, the novice was not allowed to learn because it did not have an “emergency” mode that would learn how to act in times of danger. If such mode was implemented, we recommend that strict control and mediation algorithms should be exercised, as to minimize the influence of commands that are not evaluated as similar, but are taken into consideration due to trailing effects of eventual similar controls executed earlier in the crisis.

The agents were created in a way that they would issue a control even if the Fuzzy

Associative Memory were not completely filled. As long as an output could be generated, then a command would be issued. This allows for true on-line learning while trying to operate an agent. This also shows the power of Fuzzy Mediation, as it allows a controller with an incomplete set of rules to successfully guide an agent when it can, and to learn from the expert when a rule is not present.

In this chapter we will review all the experiments that were performed on the simulators, and based on the various algorithms introduced in the earlier section. First we will look at the operations of an agent with two static controllers, two controllers that are set in their ways and do not learn or adapt to each other's performance. Then we will look at the operations of an expert controller interacting with a pre-trained A.I.-based controller. The next concept that we will explore is the analysis of different mediation engines, in order to fully understand how Fuzzy Mediation can adapt to the need of a more conservative mediation, where the expert retains most of the control, or a more trusting one, with the novice controller favored in the weight of the fusion. Then we will review experiments performed with an expert controller and a completely blank novice that will learn as the agent drives around the track. Finally, we will explore two different aspects of human interaction through the Fuzzy Mediation engine. The first set of experiments will review the operations of the simple flight simulator when two distinct controllers are involved. Then we will review the performance of humans who are trying to learn how to control the simulated object by allowing their dominant hand to collaborate with their non-dominant one.

<i>Controller ID</i>	<i>Function</i>	<i>Situation</i>
E1	$Y = X$	Expert agent
N1	$Y = X^3$	Simple novice agent
N2	$f(x) = \begin{cases} Y = X & X > 0 \\ Y = 0 & X \leq 0 \end{cases}$	Relatively novice agent that turns to the right, but can't turn left. Instead it goes straight.
N3	$f(x) = \begin{cases} Y = X^3 & X > 0 \\ Y = 0 & X \leq 0 \end{cases}$	Novice agent that turns to the right, but can't turn left. Instead it goes straight.
N4	$Y = X $	Relatively novice agent that only turns right
N5	$Y = X^3 $	Novice agent that only turns right

Table 7: Equations used to power expert and novice controllers in the first experiment

5.1. Concept of Fuzzy Mediation Applied to a Two-Agent Control

The second set of experiments goes beyond the simple analysis of Fuzzy Mediation applied to the fusion of the outputs of two different functions in a static environment. In these tests we use the simple line follower environment described earlier.

In these experiments we test the behavior of different controllers (Vincenti & Trajkovski, 2007). Also in this case, the controllers are static in nature, as each controller will be powered by a preset function that, given the position of the line to follow in relation to the present heading, will compute a new heading. The functions reported in Table 7 show the driving engines we used.

In all cases the expert controller will generate a desired heading that is exactly as suggested by the sensors. If the agent is supposed to turn 15 degrees to the left, the expert controller will output a right turn of 15 degrees. The novice controllers will vary. In the case of N1, the novice will under steer on small adjustments, then will overcompensate. In other cases, the controller will be unable to turn left, such as in controllers N2 through N5. Such radical behaviors will put Fuzzy Mediation to the test in extreme conditions.

	<i>Tight Control</i>	<i>Moderate Control</i>	<i>Loose Control</i>
<i>Average mediation weight</i>	0.59	0.62	0.82
<i>Average difference between inputs</i>	2.4	2.56	4.21
<i>Average difference between expected and actual path</i>	1.73	1.9	3.74

Table 8: Agent performance for one lap using different control levels

5.1.1. Experiments with Controller N1

The first experiment we performed involves an agent navigating through the pattern using Fuzzy Mediation to blend the controls of the expert powered by controller E1 and the novice simulated by controller N1. We studied the differences in behavior of the agent based on different levels of control, as described above. The element we monitored is the value of the mediator's weight. Table 8 shows the average values for the three levels of control using the very same controllers to simulate the inputs.

We can see that a tighter control reports an average weight of the controller that is lower when compared to the other indices. This means that, when we use the sets that correspond to a tight control of the novice's inputs, the expert retains control for more sections of the pattern than in simulations performed using looser comparisons.

Table 8 also reports two more values that are important to analyze. The first is the average difference between inputs, or the expert's input and the novice's. The next element of interest is the average difference between the expected and the actual path, measured by the difference between the expert's input and the mediated output. As we can observe from these values, the difference between the expected and the actual path is always smaller than the average difference between the trainer and the trainee's inputs.

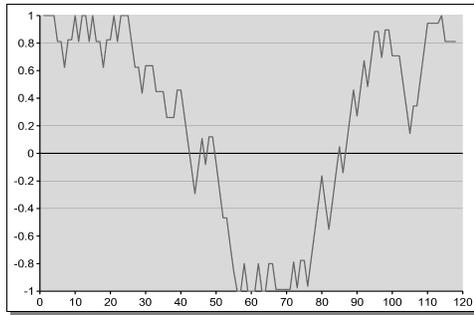


Figure 28: Shift in mediator weight during section 2, tight control

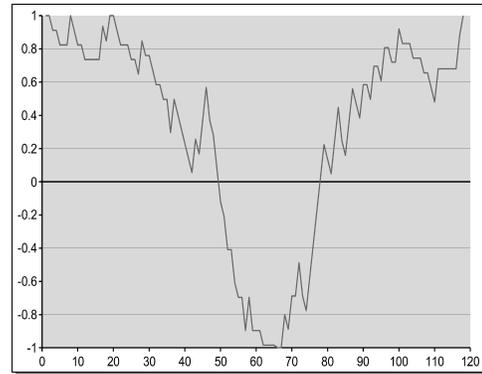


Figure 29: Shift in mediator weight during section 2, moderate control

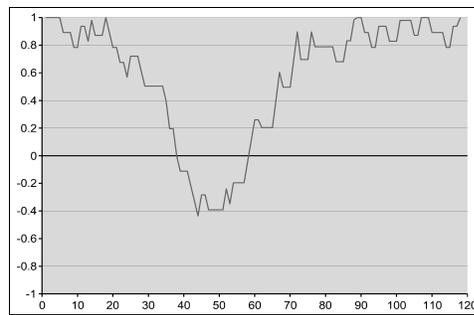


Figure 30: Shift in mediator weight during section 2, loose control

This shows that Fuzzy Mediation is successfully mediating between the two controllers, by letting the agent stay on track, and it is also reducing the error that would have been present if this algorithm was not in place.

It is important to note that the values that define the averages between the expert and the novice controllers are always different due to the fact that, with looser control, the agent finds itself further away from the track. The averages reflect the fact that the expert controller wants to over-steer in order to return to the path to follow, while the overall agent does not. In all the experimental runs we noticed that the area of the simulation field, shown in Figure 22, that consistently showed the most shift in control was number 2. Figures 28, 29 and 30 show the shift in mediator weight as the agent goes through the

	<i>Tight Control</i>	<i>Moderate Control</i>	<i>Loose Control</i>
<i>Average mediation weight</i>	0.21	0.31	0.57
<i>Average difference between inputs</i>	2.62	3.11	4.83
<i>Average difference between expected and actual path</i>	1.41	1.91	3.72

Table 9: Agent performance in section 2 of the simulation

turn. The X-axis identifies the number of steps included in the analysis, while the Y-axis identifies the range of the possible weight of the controller's mediation. It is important to remember that a mediation weight of -1 means that the expert is in control, while a weight of 1 establishes the novice as completely in control.

These figures show that, in the case of a tight control, the expert will regain significantly the control of the agent and will perform the steeper section of the turn, as shown in Figure 28.

Figure 29 instead shows that the Fuzzy Mediation that uses a moderately loose control still allows the expert to retain quite a bit of control, but overall the input of the novice is evaluated with a higher importance.

Finally, Figure 30 shows that a loose control allows the novice to take care of the majority of the control in this situation. The different nature of the controllers, as described earlier, does not allow the novice to be in control through the entire turn, as the difference in controllers' inputs are quite different.

Table 9 shows the average weight of the controller's mediator weight for the values plotted in Figures 28, 29 and 30, showing the control over section 2 of the track.

As we review the performance of the agent over this section of the track, it is important to note that, even in this case, the difference between the inputs of the two

controllers is greater than the difference of the expected and the actual output of the agent.

The following experiments were performed in order to study the interaction between an expert and a novice controller when the novice behaves quite differently from the expert. This was done in an effort to study the behavior of the Fuzzy Mediation architecture when placed in an environment where the novice controller acts in a manner that is radically different from the expert user.

5.1.2. Experiments with Controllers N2 and N3

The first simulation was performed with the expert controller based on controller E1, while controller N1 powered the novice. In the case of a turn to the right the agent showed no problems. We recorded that, in the case of a novice controller N2, the agent was leaving the pattern completely in the case of a slight turn to the left. In the case of a sharp turn to the left the agent showed some problems at the beginning of the turn, but no problem after that. This is probably due to the fact that, in the case of a slight deviation control stays unaltered, thus leaving control for a longer period of time to the novice. The second setting involved a novice controller powered by controller N3. This simulation performed very similarly to the previous one, since both simulations for the novice controller show the same response to a left turn.

5.1.3. Experiments with Controllers N4 and N5

The third set of experiments focused on controllers that behaved completely different in the case of a left turn. In order to simulate this situation we used controllers

N4 and N5 to simulate the novice controllers. We were unable to record extensive data for these experiments because the agent was not able to complete a full loop of the pattern. Our observations show that in the case of controller N4 the agent would initially follow the line, but at the first left turn it would lose control, turn completely around and then get stuck looping around itself. The same behavior was observed when controller N5 powered the novice. We then increased the value by which the weight value of the control is shifted from 0.2 to 0.5. At this point the agent with the novice controller with controller 6 performed almost a full trip around the pattern. It showed problems when it was presented with a slight turn to the left. At that point it would also start looping around itself.

5.1.4. Recommendations

The findings reported in the experiments that used novices powered by controllers N2 through N5 indicate the need to explore thoroughly Fuzzy Mediation environments that use different increments in the shift of mediation. We performed several other experiments that dealt with the shift in control, not included in this work.

5.2. Fuzzy Mediation with Static FAM

At this point we explored quite well the situations where Fuzzy Mediation interacts with static controllers in both simple settings and also settings applied to a line follower. In this next section we will place Fuzzy Mediation in an environment where it is bound to interact between an expert controller still powered by a static function, and a novice controller that relies on a pre-trained Wang-Mendel controller. In these experiments we

do not apply Fuzzy Mediation to the environment of the simple line follower, but once again we observe the overall output of Fuzzy Mediation as the two controllers feed it Y values computed using different functions.

In experiments that were reported in Section 4.2, we studied carefully the behavior of the Wang-Mendel agent in many conditions. The agent seemed to be able to learn the function $Y = X$ without problems with just a few examples. Moreover, it performed extremely well no matter what granularity the agent had. As far as the novice controller goes, we used a pre-trained Fuzzy Associative Memory. This means that the novice was shown the behavior of the function $Y = \sin(X)$ prior to commencing the experiments, and it was only allowed to perform, without improving its set of rules.

It is important to also be able to simulate different levels of agents in this pre-trained environment. The set of simulations discussed in the earlier chapter show that a situation where an agent is trained using the “Greedy” algorithm can be compared to a novice who is learning just enough about an environment to navigate through it, making some mistakes. As the novice learns more, thus observes more examples about the environment, we can relate to the results found in the agent trained using a “Random” algorithm. Finally, as a novice learns more about the environment and many of the original rules are replaced with others with higher importance, the novice becomes an expert. This situation is depicted in the simulations where the agent is trained using a “Core” algorithm. Moreover, we can assume that a novice’s understanding of a specific domain is somewhat broad at first, thus reflecting to an agent with a low granularity for the input and output domains. As the agent goes from novice to expert, we can only

<i>WL</i>	<i>SL</i>	<i>S</i>	<i>SR</i>	<i>WR</i>
$(-\infty, -\infty, -4, -3)$	$(-4, -3, -2, -1)$	$(-2, -1, 1, 2)$	$(1, 2, 3, 4)$	$(3, 4, +\infty, +\infty)$

Table 10: Boundaries of fuzzy sets used to classify the deviation between controllers

assume that the understanding of the environment increases, with an increase in the granularity of the input and output ranges. This can be simulated by having an agent perform simulations first using a low number of sets, and gradually increasing the number of sets to reach the highest possible foreseen in the environment created for these experiments.

This scenario will represent the sequence of events that we will simulate in this final round of experiments. When we introduced the concept of Fuzzy Mediation, we also addressed the analysis of the difference between the input of the trainer and the one of the trainee. For this simulation the difference between the inputs was normalized to $[-10, 10]$, the range internal to the mediation engine, and then the range was broken down into the sets reported in Table 10, which also specifies the boundaries of each set.

The fuzzy sets used map a value to a membership of 1 when the value falls between the two inner boundaries. If a value falls between an inner boundary and outer boundary on the same side, then it will have a partial membership to the set in question in according to the membership equation of choice. In this case, the equation chosen is a simple line equation. If a value falls outside of the outer boundaries, then the membership to the set in question is 0. For these simulations, the maximum absolute difference between the trainer's input and the trainee's is of one unit that is then mapped to the domain described above.

In the following simulations we will use as an index the average value of the

<i>Sets</i>	<i>Average Mediator Weight</i>
7	-0.517277228
11	0.603871287
21	0.945544554

Table 11: Core training for the sine environment with three different number of sets

mediator weight. As we described above, the control of an ideal object must be mediated between two controllers, a trainer and a trainee. The average value for this index shows which controller has most of the weight. A weight between -1 and 0 shows that the trainer has more control, with total control to the trainer if the weight is -1 . If the weight is between 0 and 1 , then the trainee has the majority of control, with control completely assigned to the trainee if the value is 1 . If the value is 0 then the control is shared equally between the trainer and the trainee.

The first aspect that we will discuss about the simulations performed is the influence of the number of fuzzy sets used for the input/output domain by the agents. Table 11 shows that, as the agent identifies the range in question with more fuzzy sets, also the weight of control shifts from a considerable average ownership of control to the trainer, as reported by the simulation with an agent based on seven fuzzy sets, to an almost complete ownership of control in favor of the trainee, as shown by the interaction with the agent based on 21 fuzzy sets.

Figures 31 through 36 show graphs reporting the input values of the trainer and the trainee as well as the mediated output. Next to each simulation we also show graphs of the value of the mediator weight after each pair of trainer/trainee input are evaluated by the Fuzzy Mediation system. The graphs correspond to the simulations reported in Table 11. These figures show how dynamic this process of mediation between two controllers

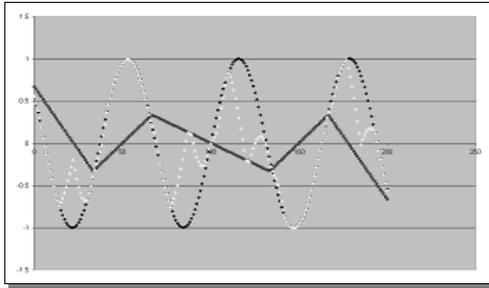


Figure 31: Input and output values for a 7-set agent, core training

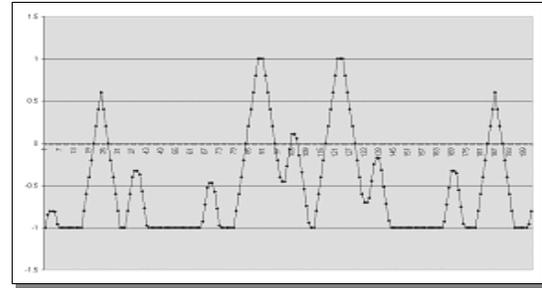


Figure 32: Control mediation value for 7-set agent, core training

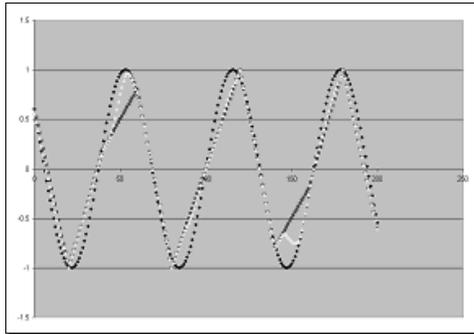


Figure 33: Input and output values for an 11-set agent, core training

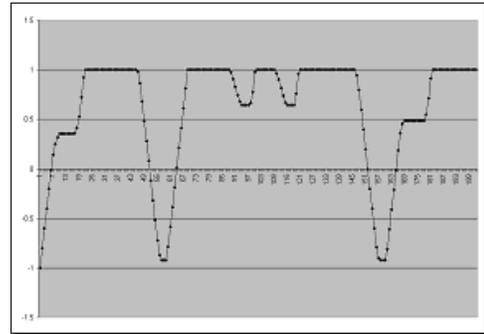


Figure 34: Control mediation value for 11-set agent, core training

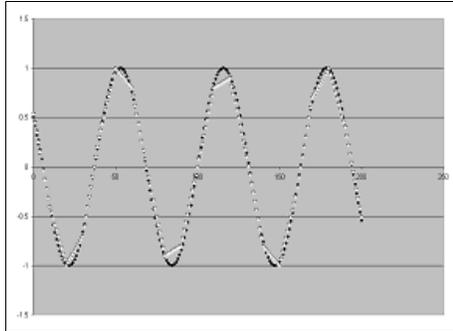


Figure 35: Input and output values for a 21-set agent, core training

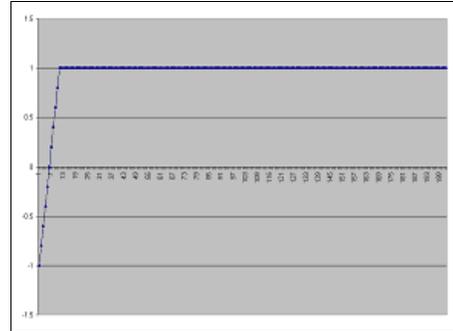


Figure 36: Control mediation value for 21-set agent, core training

can be. In Figure 31 we can see the mediation in the case where the novice can not reproduce closely the pattern shown by the expert controller. Figure 32 shows the dynamic shifting of control between the expert controller (-1) and the novice one (1). As the output of the novice does not resemble the one of the expert, then control fluctuates greatly and is mainly kept in the negative values.

<i>Training type</i>	<i>Average mediator weight</i>
Greedy	0.149693069
Random	0.110108911
Core	0.603871287
Left-to-Right	0.300168317

Table 12: Avg. mediator weight for 11-set agent (novice), different training algorithms

Figure 33 shows the pattern created by a novice controller that can reproduce the control shown by the expert. In this case the two outputs are closer, showing that, although not completely in control, the weight shifts more towards the novice, as shown in Figure 34. Finally, Figure 35 shows the event in which the novice can reproduce very accurately the behavior of the expert, thus retaining control the whole time, as we can see from Figure 36.

We also analyzed the performance of several training algorithms and the average weight of the mediator, reported in Table 12. For this particular example we chose an agent set with 11, and we can see that the agent that received “Core” training showed the best performance, keeping the average mediator value well over the side of the trainee for matters of control.

5.3. Mediation Algorithm Experiments

An important component of Fuzzy Mediation is the engine that drives the regulation of the mediation between inputs. The experiments that have been analyzed so far used a linear function, but we can replace such component with any other equation. The main objective for this experiment was to analyze the behavior of the Fuzzy Mediation engine

as we use different mediation equations.

Changing the equation that drives the mediation engine is an important step. When we use a linear function we can see that control shifts from one controller to the other in a linear manner. When control is supposed to be for 75% in the hands of the novice and 25% in the hands of the expert, then that's how the overall output is computed. This situation is not ideal in all cases though. Sometimes the situation may require that the novice be given less control, unless the outputs of the two controllers are extremely similar. For that reason we may use a function that keeps the overall weight of the expert more dominant over the one of the novice until the novice is supposed to be at least 85% in control. At that point the weighed output of the trainee will actually affect more than 50% of the overall output of Fuzzy Mediation. This can be achieved by switching the function that powers the mediation engine. In these experiments we will test several equations that will affect the weight of each of the outputs to be used to calculate the final value of the mediation. The equations tested are the following:

$$Y = X \quad \text{Equation M1}$$

$$Y = \log(X) \quad \text{Equation M2}$$

$$Y = \exp(X) \quad \text{Equation M3}$$

$$Y = X^3 \quad \text{Equation M4}$$

$$Y = \sqrt[3]{X} \quad \text{Equation M5}$$

It is important to note that the output of equations M2 and M3 was normalized to the range of $[-1, 1]$, while the others did not need normalization, as the results fell within this range. In order to be able to visualize the relation that these equations create within

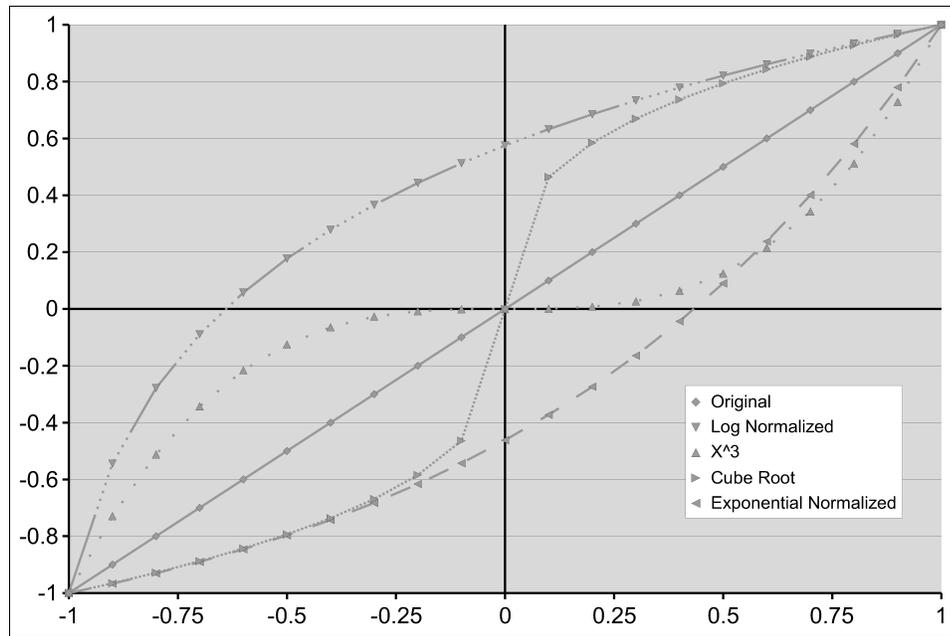


Figure 37: Mediation engines

the mediation engine, we can review Figure 37. These equations will determine the dynamics as control shifts in favor of either agent.

For this experiment we used the simple two-controller line follower. The two agents were represented by an expert agent powered by the equation $Y = X$ and a novice controller powered by $Y = X^3$.

The experiment reported in this section also explored the dynamics of the agent with two controllers with different increments within the process of mediation. Going back to the description of the algorithm, we explained how control shifts towards the trainer or towards the trainee with a certain increment over the domain $[-1, 1]$, with -1 signifying total control to the trainer and 1 allowing the trainee to control the object. For other experiments we used the increment of 0.2, which is a choice that allows a dynamic shift of control that is not either too rapid or too slow.

	<i>Tight</i>	<i>Moderate</i>	<i>Loose</i>
Lap 1	0.591630435	0.621014947	0.82562658
Lap 2	0.587499453	0.625134209	0.807701818
Lap 3	0.599046749	0.629752809	0.802773797
Lap 4	0.592637091	0.625473723	0.810562727
Lap 5	0.600289971	0.621890988	0.809399855
Lap 6	0.594071299	0.623954958	0.830268743
Lap 7	0.5858	0.625674231	0.812529946
Lap 8	0.58594359	0.630120029	0.825644846
Lap 9	0.603001442	0.618650759	0.814850542
Lap 10	0.590514801	0.6208102	0.817715729

Table 13: Agent performance over 10 laps

Before we get any further into the description of the experiments, it is important to note that each run was conducted over a single lap around the course. While conducting preliminary experiments, we noticed that the performance of an agent for each lap is quite similar through a 10-lap test. The values shown in Table 13 show the average mediation weight for different tightness of control after each lap.

We analyzed the average and the standard deviation of the values reported in Table 13, shown in Table 14. We can see that the standard deviation for each lap is small enough for which a lap performed under a tight comparison method cannot be misinterpreted as a lap that was mediated through a moderate controller. When we analyze the performance of a moderate tightness of control versus a loose one, we can see that the difference is large enough for which confusion is very unlikely.

	<i>Tight</i>	<i>Moderate</i>	<i>Loose</i>
Average	0.593043483	0.624247685	0.815707458
Standard Deviation	0.006063892	0.003776346	0.008954575

Table 14: Average and Standard Deviation values for the base experiment

	<i>0.05</i>	<i>0.1</i>	<i>0.2</i>	<i>0.5</i>
Tight	0.598187589	0.591931181	0.593774938	0.609667733
Moderate	0.654717268	0.630193206	0.622115398	0.647160142
Loose	0.837450808	0.807863588	0.823889831	0.810044263

Table 15: Performance of the M1 mediation engine

	<i>0.05</i>	<i>0.1</i>	<i>0.2</i>	<i>0.5</i>
Tight	0.278387154	0.246983057	0.279714387	0.326583895
Moderate	0.319678584	0.329285561	0.310036209	0.383241661
Loose	n/a	0.645599929	0.597444288	0.624077862

Table 16: Performance of the M2 mediation engine

	<i>0.05</i>	<i>0.1</i>	<i>0.2</i>	<i>0.5</i>
Tight	0.729839957	0.737176849	0.727778653	0.759396608
Moderate	0.780996076	0.774061232	0.769685714	0.770878451
Loose	0.889485048	0.896666667	0.89193465	0.883299393

Table 17: Performance of the M3 mediation engine

	<i>0.05</i>	<i>0.1</i>	<i>0.2</i>	<i>0.5</i>
Tight	0.701056948	0.698215756	0.703043571	0.690017806
Moderate	0.773597047	0.749746795	0.730972973	0.73344274
Loose	0.902675574	0.909070163	0.888818697	0.863217822

Table 18: Performance of the M4 mediation engine

	<i>0.05</i>	<i>0.1</i>	<i>0.2</i>	<i>0.5</i>
Tight	-0.728596407	-0.626837683	-0.579029506	-0.551495028
Moderate	-0.70347103	-0.604550662	-0.561449788	-0.52778709
Loose	-0.302748483	-0.221367281	-0.207888651	-0.186754825

Table 19: Performance of the M5 mediation engine

The equations were chosen with specific purposes in mind. We wanted to analyze the behavior of each one so that we could match a certain purpose for the mediation to each engine. Some equations leave the control to the trainer for a longer time, while

others favor the trainee. Possible applications of this algorithm may require certain settings over others, and this work outlines some suggestions for the initial setup of the environment. Tables 15 through 19 show the average mediation weights for the simulations just described.

We can see that the mediation engine M1 offers a good balance among all the ones tested. The average mediator weight for any of the simulations performed, compared to the same setup of other equations always shows to be in the middle of the range. Engine M2 shows a control that favors the trainer when performing the simulations. Engine M3 instead offers the highest average mediation weight, showing that the trainee was the controller that was most in control of the agent. Engine M4 also shows a preference for the trainee, but not as much as Engine M3. Finally, Engine M5 seems to leave more control to the trainer, making this the simulation that leaves the least leeway to the trainee.

The second aspect that is relevant to this work is the measure by which control was shifted towards one or the other controller. As stated previously, the default value is an increment of 0.2. In these simulations we tested the increments of 0.05, 0.1, 0.2 and 0.5.

As we review the results of the experiments, we can see that in the majority of the cases the different increment in shift of control does not play a significant difference. Although values fluctuate up and down by a few points, generally the tightness of control does not play a major factor in preferring the trainer versus the trainee.

The first point that is interesting is that the simulation with the increment set to 0.05, a loose control and a mediation engine driven by Engine M2 did not lead to a successful completion of a lap. The agent was not able to perform a full rotation. Although this

mediation engine is considered to prefer the trainer to the trainee, an increment of 0.05 seems to be too slow to let the trainer gain back control in situations where the performance of the trainee is not adequate.

The most notable effect of the different increment in the shift of the mediation can be seen in the simulations driven by Engine M5, the one that retains control closest to the trainer. As the increments decrease, the trainer retains control even more through the simulation.

When we first explored possible equations that would power the mediation engine, we assumed that we were to base our decision of which would relinquish more or less control dynamically by analyzing the behavior of the equation along the range $[-1, 1]$. The data shows that, given the simulation setup described earlier, the most important interval to keep in mind is actually the behavior within the range $[-0.5, 0.5]$, as that seems to be the most dynamic portion of control, and at the same time the range where we can find the mediation the majority of the time.

Overall, we feel that we can choose which area needs to be most dynamic as far as the mediation goes by selecting the appropriate mediation engine. The findings of these experiments suggest that Engine M1 offers a balanced mediation between the two controllers.

Engine M3 and M4 are the ones that leave the most room to novice controllers, thus stimulating the exploratory nature of certain learning algorithms that construct rules as the agent explores and creates associations that show good behaviors and bad behaviors.

Engine M2 and M5 are preferable in situations where the learning mechanism of the novice agent thrives on examples that show it what to do, instead of also showing what not to do. These equations also seem fitted for environments where there is a need for an agent to learn rules that are very similar, if not identical, to the expert's knowledge.

Moreover, a smaller value regulating the shift of control between the trainer and the trainee seems more appropriate for precision type agents. We do recommend though that the smallest increment adopted for situations similar to the one described by these experiments be at most reduced to 0.1, as a smaller value may lead to problems in the performance of the task.

5.4. Fuzzy Mediation with a Dynamic FAM

In the following set of experiments we will analyze how Fuzzy Mediation affects the operations of an agent that is controlled by a static expert, powered by $Y = X^3$ and a novice powered by a dynamic Wang-Mendel Fuzzy Associative Memory that starts without any knowledge. In these experiments we will use the environment of the simple line follower. Since we have not tested the behavior of the Wang-Mendel algorithm with this particular function, we will first perform some tests to establish a baseline for the behavior to expect. Moreover, we never tested this particular controller applied to our particular simulation environment, so this will give us a chance.

5.4.1. Preliminary Study

Before we could go on with testing Fuzzy Mediation in agents with two controllers placed in a live situation, it was necessary to study the behavior of the agent based on the

<i>FAM sets</i>	<i>Training pairs required</i>
3	9
5	19
7	25
9	25
11	25
13	38
15	39
17	45
19	71
21	103

Table 20: Training pairs required for greedy training, different granularity agents

Wang-Mendel algorithm. To do so, we reviewed the many training environments described earlier in this work and put them to the test. Each agent was to be trained and then placed in the simulated environment. The training was based on the equation $Y = X^3$. The agent then had to perform five laps, and we recorded the expected output and compared it to the actual output. Before we analyze these preliminary experiments any further, it is important to note that, given a greedy training environment, the agent needs a different number of training pairs to become minimally functional, in the sense that its FAM contains one association for each location. The number of minimum training examples is given in Table 20. The X-value of each training pair was picked at random.

Some of the agents with different granularities require the same amount of training pairs, or a very similar number, because the values are random. In some decompositions the random values may teach redundant rules. As the width of each set decreases with a finer granularity, the redundant examples may actually cover the span of two or three rules, thus keeping the number of training values at similar levels, but still training the agent fully.

	<i>125</i>	<i>250</i>	<i>500</i>
3	31.98654679	31.98696547	31.98687252
5	3.612871604	3.61294211	3.612904595
7	1.502545679	1.50320008	1.50265563
9	0.725142265	0.725153279	0.725158206
11	2.859757173	2.859788277	2.859331915
13	1.508049472	1.508250578	1.508573294
15	0.91887035	0.918852038	0.91903639
17	2.143460895	2.143728468	2.143307533
19	1.521651934	1.520235104	1.521538381
21	0.923157012	0.923287203	0.923099485

Table 21: Avg. abs. diff. between expected/actual output, different count training pairs

As the simulations that we will perform will be based on random training pairs shown to the agents, it was necessary to review how the behavior of the agent changed as the number of examples shown increased. Table 21 shows the changes in the absolute difference between the expected and the actual output. We can see that as the number of training examples increased, the average error was quite stable. There was no noticeable difference between laps performed with a prior training of 125, 250 or 500 examples. It is important to note that the high difference in the agents with three fuzzy sets was determined by the fact that they responded by making sharp turns to the left and to the

	<i>Greedy</i>	<i>Random</i>	<i>Core</i>	<i>Left-to-Right</i>
3	n/a	31.98687252	31.98638183	31.98658987
5	3.6130439	3.612904595	3.613663374	3.612957563
7	1.501601203	1.50265563	1.50229543	1.502666797
9	0.725262943	0.725158206	2.377974771	0.725014606
11	2.79416703	2.859331915	2.859455812	2.859402731
13	1.507901291	1.508573294	1.508275538	1.507569353
15	2.676876176	0.91903639	0.919025033	0.919022209
17	2.144023474	2.143307533	2.144346614	2.144091617
19	1.521634656	1.521538381	1.5215416	1.522203731
21	0.923343914	0.923099485	0.923055746	0.923114618

Table 22: Avg. abs. difference between expected and actual output, different trainings

right, thus performing extreme motions. In real life, a robotic agent powered by a 3-set Wang-Mendel based controller could not function.

Table 22 shows the absolute average difference between expected and actual output for the various training situations. Just like in the previous case, an agent with 3 sets completed the course in a manner that would have been extremely strenuous for a real agent to withstand. In the case of greedy training with three sets, the agent actually did not complete even one lap of the simulation. It is important to note that these seems to be a trend in performance when we analyze the random training runs, where the performance improves steadily for agents with three sets, all the way to agents with seven sets. At that point though the average error increases again, to then keep decreasing until the agent has 15 sets. An agent with 17 sets then performs worse than the previous agent. The trend then follows the same pattern, with the error decreasing until the last agent, the one with 21 sets, declines in error. These results, as they are averages, create the baseline to which we can compare later runs where Fuzzy Mediation is used to share control between trainers and trainees.

5.4.2. Observation of a Single Lap

An agent placed in an environment where it is supposed to start functioning immediately is similar to an agent that functions right out of greedy training. The major difference between the training environment tested to get the baseline of the behavior and the agents placed in the following simulations is that the greedy training shown earlier assumes that the entire rule base needs to be filled, while the agents that we will use may potentially become operative after being exposed to just one training set. This means that

the agents will be performing on a sub-minimum greedy type training to start with, and they may complete the simulation with the set of associations not complete for the whole domain, because they were not exposed to all possible examples.

A sub-minimum greedy type training refers to the situation in which an agent creates a rule based on a single training pair. Given the single association present in a FAM with more than a single set, the agent still requires many associations to be filled in order to become minimally functional. In a sub-minimally trained agent, should an event relatively similar to the one that generated the association stored be present, then the agent will be able to react. Given any other scenario, the agent will not have any associations to reference in order to generate an output.

Now that we have a good understanding of how an agent based on the Wang-Mendel algorithm behaves on the simulation created for these experiments, we can start analyzing the performance of the Fuzzy Mediation engine. As the agent is placed in the environment with no knowledge of anything, the fuzzy associative memory starts empty. As the agent collaborates and learns from the expert, the rule base will start filling up. As the rule base fills up, the agent will also be able to start generating values to control the object. When the agent cannot find a rule for a certain situation, then it sits back and simply observes and learns. Otherwise, it will generate its own output, then also observe the behavior of the expert and learn. This is done because we do not assume that, once a rule is set, that is the best possible rule that may be generated. We are trying to test completely dynamic systems that keep learning as the operations continue.

	Step 0	Step 1	Step 2	Step 4	Step 22	Step 24	Step 25	Step 34
Loc 0	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)
Loc 1	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)
Loc 2	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)
Loc 3	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)
Loc 4	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)
Loc 5	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)
Loc 6	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)
Loc 7	-1 (0)	-1 (0)	-1 (0)	7 (0.92)	7 (0.92)	7 (0.92)	7 (0.92)	7 (0.92)
Loc 8	-1 (0)	-1 (0)	7 (0.8514)	7 (0.8514)	7 (0.8514)	7 (0.8514)	7 (0.8514)	7 (0.9114)
Loc 9	-1 (0)	-1 (0)	-1 (0)	-1 (0)	7 (0.639)	7 (0.6952)	7 (0.799)	7 (0.799)
Loc 10	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)
Loc 11	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)
Loc 12	-1 (0)	9 (0.6068)	9 (0.6068)	9 (0.6068)	9 (0.6068)	9 (0.6068)	9 (0.6068)	9 (0.6068)
Loc 13	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)
Loc 14	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)	-1 (0)

Table 23: Evolution of the FAM as the simulation proceeds

The first few steps of the agent are crucial. These are the steps when the first associations start forming, old rules get replaced quickly, and perhaps the novice controller skips a few steps because of a lack of rules. A step is defined as follows:

- the robot pauses for a very short time;
- it assesses its position in relation to the line that is to be followed;
- it receives values from the expert and, if possible, from the novice, as to how the heading should change to follow the line;
- it evaluates the two inputs without Fuzzy Mediation if only one input is present, otherwise it uses Fuzzy Mediation to generate a single value;
- it takes a step in the direction computed.

Table 23 shows the evolution of a 15-set FAM within the Wang-Mendel controller during the first steps of an agent not previously trained. Only steps where changes are

<i>Memory location</i>	<i>Value</i>	<i>Association Weight</i>
0	-1	0
1	-1	0
2	-1	0
3	6	0.44
4	6	0.5529
5	7	0.8134
6	7	0.9702
7	7	1
8	7	0.9702
9	7	0.8134
10	8	0.5429
11	-1	0
12	9	0.6068
13	-1	0
14	-1	0

Table 24: Final FAM

made are reported.

As expected, during the first step the FAM is empty. An empty location is indicated by a value of -1 with rule strength of 0. After the first step, the first rule is created. Step two also sees a new rule placed in memory. We can observe the first substitution of a rule in step 24, where the association stored in location 9, previously a 7 with rule strength 0.639 is replaced with another rule, still pointing to the output domain at location 7, but this time with a strength of 0.6952. This shows the extreme dynamic nature of the approach. The FAM after the first lap is shown in Table 24.

As mentioned earlier, the first few steps of the simulation are extremely important. Figure 38 shows the outputs of the trainer and of the trainee (if one is present) as well as the mediated output.

We can see that the output of the agent for the first 100 steps is still rather raw. The

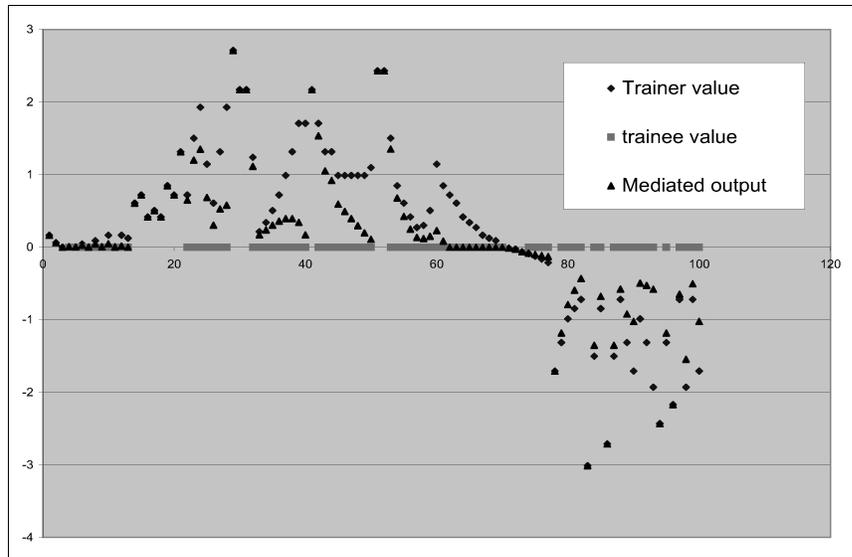


Figure 38: Single Lap Observation - Raw data between steps 1 and 100

expert maneuvers the agent with right and left turns, while the primitive rule base of the agent keeps it going on a straight line. It is important to note that Fuzzy Mediation still acknowledges the novice's input, but manages to keep the agent on the line. This is particularly evident between steps 30 and 40, where the mediated output seems to lean towards the novice, but at some point, as the novice doesn't have a rule to offer, control

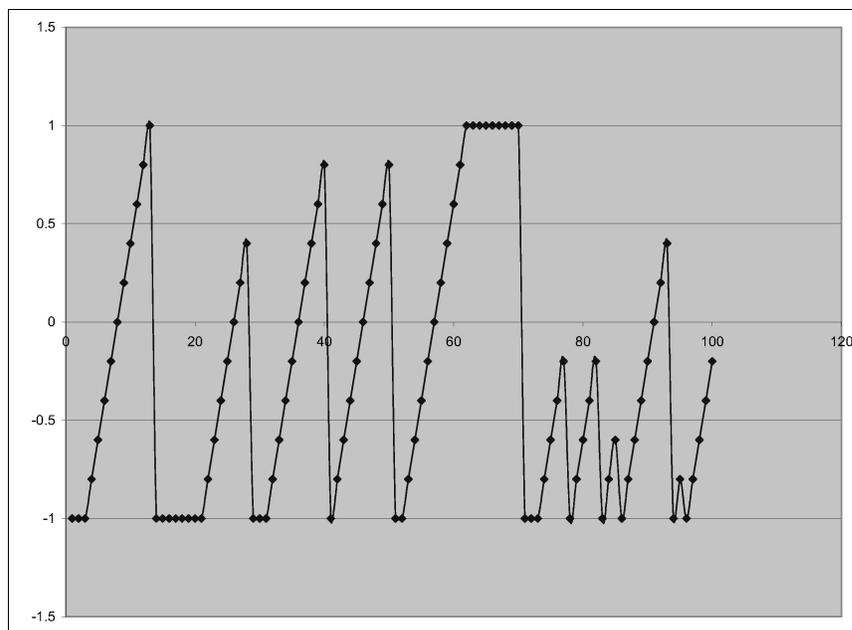


Figure 39: Shift in control between steps 1 and 100

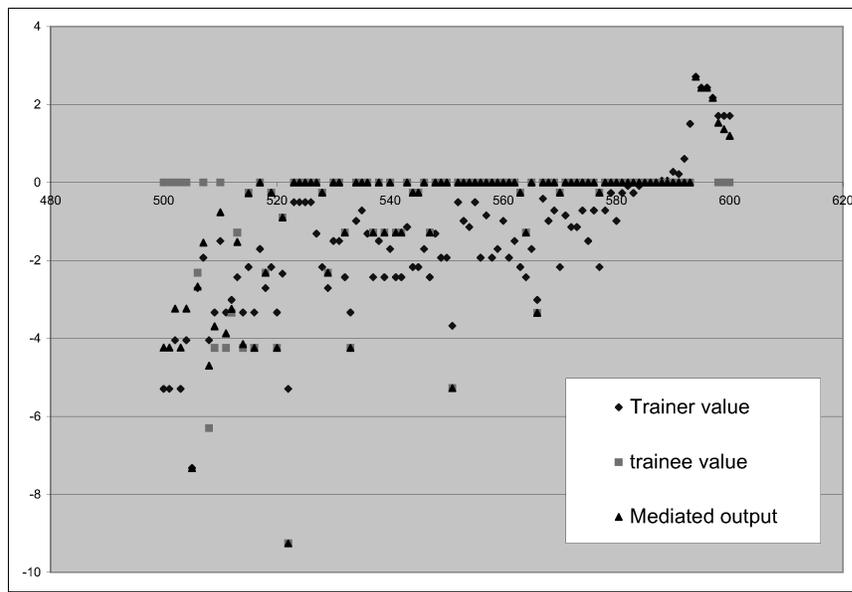


Figure 40: Raw data between steps 500 and 600

goes immediately to the expert. A missing rule for the novice is shown by a break on the squares along the X-axis.

Figure 39 shows the value of the mediation weight for the very first 100 steps, to follow along with Figure 38. We can see that the majority of the time the weight is below 0, which shows that the expert is predominant in controlling the agent. Sharp changes in control from the novice to the trainer denote a lack of output coming from the trainee. When this happens, weight is shifted immediately back to the trainer without the possibility of making up the rule once the observation of the expert has taken place.

Figure 40 shows the interaction between the expert and the novice through the Fuzzy Mediation engine between steps 500 and 600. We can now see that the novice is a bit more dynamic. Also the mediation values are more dynamic and follow the novice more than the agent. It is important to note that the difference between novice and expert in the majority of the cases is about one to two degrees, thus not showing enough difference to revert control back to the trainer.

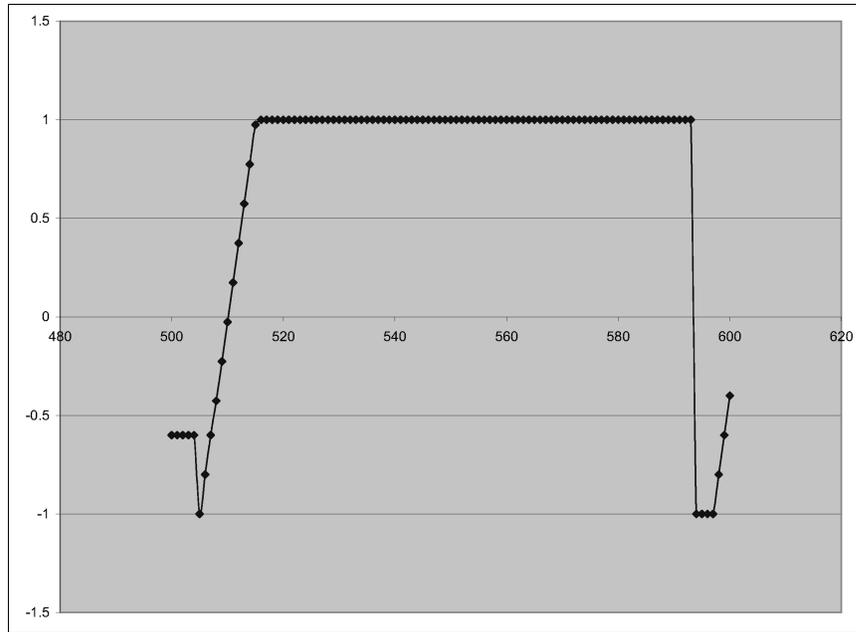


Figure 41: Shift in control between steps 500 and 600

Figure 41 shows the mediation weight for steps 500 to 600. We can notice an improvement over the first 100 steps, where the trainer mainly controlled the agent. In this case, the trainee is the controller that generates the output that will then be followed. We should notice the gradual shift in control between steps 500 and 520. Also, the rapid shift back towards step 590, then the trainee, in full control up to that moment, doesn't have a rule to generate an output, reverting control back immediately to the trainer.

5.4.3. Behavior of Different Granularity Agents

Now that we established the dynamics of a Wang-Mendel controller embedded in the Fuzzy Mediation framework alongside an expert controller, we can look at how this fusion algorithm controls and oversees the actions of the dual-controller agent. The following tests were conducted by allowing the agent to perform five laps around the circuit. The values reported in the following tables are averages. We also report the

	<i>0.05</i>	<i>0.1</i>	<i>0.2</i>	<i>0.5</i>
3	27.79409843	1.312529281	1.402325689	1.352455377
5	1.34247985	1.391489871	1.621951231	1.52001326
7	1.496585557	1.44884264	1.400326137	1.397458119
9	0.862083528	0.869910574	1.227982307	1.262718148
11	1.581845732	1.515014865	1.480653772	1.432990652
13	1.493120702	1.454657638	1.391637812	1.381869954
15	1.03452667	1.043668079	1.083786812	1.089572747
17	1.468120279	1.440023594	1.427739556	1.437028694
19	1.43999095	1.425612661	1.395166471	1.385289127
21	1.182239275	1.202574921	1.213496064	1.236640953

Table 25: Average absolute difference between inputs, Trainer/Trainee

relative standard deviations to better relay the idea of consistency of control through the simulations.

For each set of experiments we review three indices. The first is the average absolute difference between the trainer's and the trainee's input. The second set of values we review is the average absolute difference between the trainer's input and the mediated output. This shows the gap between the expected path and the actual one. The last index shows the average mediation weight through the simulation. This value gives us an idea of where control was, on average, through the simulation.

The first set of experiments was conducted to study the behavior of an agent based on the Fuzzy Mediation framework with different increments in the shift of control. We studied four increments, 0.05, 0.1, 0.2 and 0.5. Tables 25 through 30 report the values associated with this experiment. The following tests were conducted using tight control and a linear mediation engine.

	<i>0.05</i>	<i>0.1</i>	<i>0.2</i>	<i>0.5</i>
3	13.12348998	0.092561504	0.286043738	0.209784446
5	0.158695073	0.287924219	0.21596072	0.09127331
7	0.096711588	0.060945197	0.028263931	0.014436151
9	0.145107938	0.156009756	0.011952529	0.111075995
11	0.069689146	0.038471348	0.023988391	0.012988763
13	0.108463523	0.074681005	0.053091526	0.014479518
15	0.044154898	0.030223696	0.032399062	0.036106306
17	0.153798584	0.069298853	0.047373052	0.01233741
19	0.057460595	0.039065806	0.037616835	0.008045747
21	0.027989628	0.014087175	0.015076035	0.014086559

Table 26: Standard deviations associated to values in Table 25

In Table 25 we can observe a trend that we already noticed earlier, as we were analyzing the behavior of the Wang-Mendel algorithm. This goes to show that the basic behavior goes unchanged even when Fuzzy Mediation is in place. We see that the average difference between inputs decreases in controllers as the number of sets increases from 3 to 9, to jump back up with 11 sets, then to fall again, and with another increase at 17 that then reports a decrease up to 21 sets.

Table 27 reinforces the fact that Fuzzy Mediation operates effectively in controlling

	<i>0.05</i>	<i>0.1</i>	<i>0.2</i>	<i>0.5</i>
3	27.6910644	0.75281133	0.657587963	0.643396184
5	0.925552593	0.911600296	0.850702545	0.822801157
7	1.165573821	1.077729376	0.949320318	0.869891376
9	0.744441857	0.743120531	0.762272108	0.715091463
11	0.99851092	0.902539006	0.827331113	0.785240844
13	1.178295185	1.060705074	0.959159162	0.882409492
15	0.887439541	0.858021418	0.929235836	0.929152322
17	0.93273616	0.825232408	0.801905371	0.791130261
19	1.074184448	1.004686385	0.957506805	0.891701756
21	1.084683656	1.080608024	1.116970481	1.115968865

Table 27: Average absolute difference for output values, Trainer/Mediated Output

	<i>0.05</i>	<i>0.1</i>	<i>0.2</i>	<i>0.5</i>
3	13.35388049	0.170156799	0.166075739	0.171977521
5	0.025397393	0.005942634	0.011549945	0.021669108
7	0.19613108	0.125547426	0.08553906	0.04316758
9	0.105802319	0.092294427	0.096699297	0.078535333
11	0.087047932	0.044839226	0.027569965	0.011344878
13	0.214094697	0.152610495	0.115289845	0.07240453
15	0.197209075	0.171910634	0.068979832	0.14171746
17	0.222833636	0.110057696	0.06356347	0.018949761
19	0.118639117	0.091115735	0.098963035	0.03310534
21	0.182013914	0.118423704	0.09845797	0.07451725

Table 28: Standard deviations associated to values in Table 27

the agent successfully. When compared to the preliminary tests, the difference between expected and actual paths is much smaller with Fuzzy Mediation.

Table 29 reports the average mediation weight through the simulations. We don't see a net difference among mediation increments as far as average values go. We do find smaller standard deviations through the runs, pointing to the fact that the behaviors of the agents become more similar to the behavior of the expert controller.

	<i>0.05</i>	<i>0.1</i>	<i>0.2</i>	<i>0.5</i>
3	0.80019739	0.253631085	0.051753052	0.087371515
5	0.595154921	0.672431005	0.359728912	0.370986689
7	0.610873924	0.571222238	0.492531268	0.433304635
9	0.824027093	0.828644814	0.494606714	0.448901035
11	0.361765277	0.350665349	0.333094972	0.35109616
13	0.611174833	0.539784438	0.502858064	0.451702341
15	0.760492116	0.734649705	0.818732258	0.828046036
17	0.361418014	0.316927634	0.333975344	0.349390053
19	0.522060794	0.488918727	0.489181452	0.449883774
21	0.847817763	0.831882333	0.879761249	0.861320237

Table 29: Average mediator weight

	<i>0.05</i>	<i>0.1</i>	<i>0.2</i>	<i>0.5</i>
3	0.446772217	0.412653959	0.376895729	0.370193358
5	0.181510685	0.077913472	0.045471403	0.031134651
7	0.163223901	0.092101668	0.062863737	0.026703812
9	0.382576458	0.354902126	0.167249124	0.139269571
11	0.072604653	0.029691769	0.021820949	0.015485815
13	0.185215957	0.12043154	0.082918228	0.061891446
15	0.282361964	0.241806482	0.106275254	0.207252578
17	0.165669507	0.08191586	0.044816938	0.016677059
19	0.12062168	0.092404749	0.096250695	0.036288998
21	0.25139668	0.156052637	0.139221	0.100572417

Table 30: Standard deviations associated to values in Table 29

The aim of Fuzzy Mediation is not to hide the true nature of the learning algorithm, but it intends to create a successful mediation between an expert controller and a novice one. Given the fact that the difference between the expected path and the actual path of these simulations is smaller than the difference between the expected and the actual one recorded during testing, we can say that there is an improvement as far as the control of the agent goes. Fuzzy Mediation does help in minimizing the error in control, while an agent alternated between forming the knowledge base required for control and using these data to lead the agent.

It is important to note that the agent based on 3 sets performed better than the initial tests. Although the increment of 0.05 did not seem fit to this type of agent, in other situations Fuzzy Mediation was actually able to maintain the robot within reasonable operational limits, without stressing the robot nearly as much as shown in preliminary analyses. The high values for the run with 0.05 are probably to attribute to the fact that the shift between controllers is too slow to be effective.

5.4.4. Behavior with Different Tightness of Control

Now that we have a better understanding of the behavior of Fuzzy Mediation with a Wang-Mendel agent as the increment changes, we should explore different conditions as far as the comparison of inputs goes. In previous experiments we showed that a looser comparison method increments the average mediation weight, thus shifting control towards the novice. In the following experiments, we expect to observe similar findings. It is important to analyze the application of such comparison schemes in a situation where the agent is actually learning dynamically.

Tables 31 through 36 show the average values recorded during the experimental runs. The environment was set to have a linear mediation engine and a 0.2 increment for the shifting in control.

Table 31 shows mixed results in the analysis of the difference between inputs. Some controllers report a performance with a lower difference, while others seem to do worse. Also the standard deviation, reported in Table 32, doesn't show a clear improvement.

	<i>Tight</i>	<i>Moderate</i>	<i>Loose</i>
3	1.402325689	1.798516009	33.55538149
5	1.621951231	1.728410787	2.291708987
7	1.400326137	1.622328646	1.687622106
9	1.227982307	1.021026126	0.932825952
11	1.480653772	2.15435881	3.118212377
13	1.391637812	1.658798027	1.692578621
15	1.083786812	1.0882114	1.067870902
17	1.427739556	2.068937164	2.501348338
19	1.395166471	1.720571822	1.694811222
21	1.213496064	1.231320231	1.214149813

Table 31: Average absolute difference between inputs, Trainer/Trainee

	<i>Tight</i>	<i>Moderate</i>	<i>Loose</i>
3	0.286043738	0.077436919	0.267038239
5	0.21596072	0.200515851	0.126140159
7	0.028263931	0.119333004	0.026037123
9	0.011952529	0.295647587	0.297688122
11	0.023988391	0.036428263	0.514699386
13	0.053091526	0.113625838	0.062172307
15	0.032399062	0.069916849	0.029779947
17	0.047373052	0.199339145	0.383174875
19	0.037616835	0.022906118	0.022400275
21	0.015076035	0.018841286	0.028587507

Table 32: Standard deviations associated to values in Table 31

The same situation can be seen in Tables 33 and 34. The actual performance of the agent fluctuates as different controllers are in use.

Finally, Tables 35 and 36 show the average mediator weight and the relative standard deviations for these tests. As expected, the looser the comparison method is, the more the novice is in control.

The only point that needs to be addressed is that the controller with 3 sets and loose control is really a false result in the sense that, although it completed the test, an agent

	<i>Tight</i>	<i>Moderate</i>	<i>Loose</i>
3	0.657587963	1.270374966	33.55106428
5	0.850702545	1.221386845	1.979608892
7	0.949320318	1.461771295	1.648132579
9	0.762272108	0.856619983	0.886490237
11	0.827331113	1.572272439	3.060123098
13	0.959159162	1.547881115	1.627409987
15	0.929235836	1.024351002	1.020084129
17	0.801905371	1.531439651	2.375539356
19	0.957506805	1.631761227	1.658039231
21	1.116970481	1.182148938	1.199750935

Table 33: Average absolute difference between outputs, Trainer/Mediated Output

	<i>Tight</i>	<i>Moderate</i>	<i>Loose</i>
3	0.166075739	0.064634786	0.276690965
5	0.011549945	0.03726115	0.087870255
7	0.08553906	0.253609397	0.108247067
9	0.096699297	0.129043173	0.199702192
11	0.027569965	0.049069022	0.644526427
13	0.115289845	0.302434086	0.190854624
15	0.068979832	0.011969967	0.040495919
17	0.06356347	0.254724742	0.561589928
19	0.098963035	0.062155139	0.050978592
21	0.09845797	0.047758918	0.039124066

Table 34: Standard deviations associated to values in Table 33

wouldn't be able to physically stand the type of control involved with variations in heading by up to 45 degrees consistently throughout the simulations. The weight of control shows that the novice was in charge of the motions the majority of the time simply because the novice was technically applying the correct rule, as the expert would. The problem though is that the overcorrection of the heading resulted in the mirror image of the same situation on the other side.

Also in this case, Fuzzy Mediation allows the novice controller to interact with the agent immediately. Most setups will work with loose control, and all will work with

	<i>Tight</i>	<i>Moderate</i>	<i>Loose</i>
3	0.051753052	0.576059764	0.998234003
5	0.359728912	0.746021252	0.868491258
7	0.492531268	0.854279288	0.96554402
9	0.494606714	0.838799138	0.955347183
11	0.333094972	0.576373678	0.953157542
13	0.502858064	0.894283646	0.93740033
15	0.818732258	0.940643238	0.952892
17	0.333975344	0.58713215	0.901958041
19	0.489181452	0.921768908	0.967501457
21	0.879761249	0.948840218	0.983359268

Table 35: Average mediator weight

	<i>Tight</i>	<i>Moderate</i>	<i>Loose</i>
3	0.376895729	0.24892196	0.00394889
5	0.045471403	0.066736548	0.020054237
7	0.062863737	0.128199864	0.077045914
9	0.167249124	0.265345389	0.098447404
11	0.021820949	0.031135844	0.104742921
13	0.082918228	0.195283687	0.134862209
15	0.106275254	0.080278023	0.076231806
17	0.044816938	0.081678044	0.170507516
19	0.096250695	0.067224428	0.055102752
21	0.139221	0.062440695	0.037209807

Table 36: Standard deviations associated to values in Table 35

moderate and tight comparisons. In most cases, even comparing the inputs loosely allows for a better performance overall. If we compare the average difference between the trainer's input and the mediated output to the average difference between expected input and actual output in the preliminary tests, we see that Fuzzy Mediation still allows the agent to perform with a motion that is closer to the expert controller's actions.

The experiments just reviewed then suggest that novice controllers are more in control of the agent, as the comparison sets get looser. It is important though to understand that controllers that are at high risk for misguiding the agent need to be under strict supervision, in order to avoid problems to the overall operations.

5.4.5. Behavior with Different Mediation Algorithms

Finally we take a look at the behaviors of agents with different mediation engines. In earlier experiments we saw that a linear mediation offers a good middle ground for this component of the algorithm. Mediation engines based on cube root and logarithmic operators allow the trainer to be more in control, while exponential and cube functions

	<i>M1</i>	<i>M2</i>	<i>M3</i>	<i>M4</i>	<i>M5</i>
3	1.402325689	n/a	1.19902868	1.198911996	33.62318036
5	1.621951231	1.578249457	1.69242267	1.504256891	1.905956452
7	1.400326137	1.449789064	1.394330216	1.411937777	1.754445414
9	1.227982307	0.864818954	0.919486037	1.327289042	0.833280324
11	1.480653772	1.513829413	1.5547654	1.461226068	3.24518257
13	1.391637812	1.435369575	1.392488078	1.397071087	1.748073782
15	1.083786812	1.010568407	1.010572757	1.080784671	1.059250493
17	1.427739556	1.459100662	1.39751819	1.451052908	2.623631295
19	1.395166471	1.455538708	1.378033221	1.380225751	1.722145927
21	1.213496064	1.164512807	1.181293757	1.229146197	1.216312764

Table 37: Average absolute difference between input, Trainer/Trainee

shift the focus of control on the trainee. In this experiment we will revisit all of these mediation engines putting them to the test and analyzing how they perform. The comparison level for these experiments is set to tight and the increment in the shift of mediation is of 0.2.

Tables 37 and 38 show the average abstract difference between the inputs of the trainer and of the trainee. We can see that, compared to other experiments, these levels are somewhat similar, which is normal because we did not change any of the agents.

	<i>M1</i>	<i>M2</i>	<i>M3</i>	<i>M4</i>	<i>M5</i>
3	0.286043738	n/a	0.019402537	0.123700936	0.136436573
5	0.21596072	0.013403897	0.088563375	0.013784158	0.110678958
7	0.028263931	0.032265087	0.025521945	0.034851345	0.037629959
9	0.011952529	0.051131434	0.05468226	0.150221125	0.067850602
11	0.023988391	0.011606864	0.025413902	0.01208104	0.224025995
13	0.053091526	0.045090451	0.016311428	0.01319077	0.043375732
15	0.032399062	0.025027155	0.051484314	0.043222347	0.021249416
17	0.047373052	0.050046878	0.028621252	0.018640904	0.031564577
19	0.037616835	0.043695746	0.015229204	0.023596213	0.035057977
21	0.015076035	0.05471152	0.037944109	0.009377149	0.036589646

Table 38: Standard deviations associated to values in Table 37

	<i>M1</i>	<i>M2</i>	<i>M3</i>	<i>M4</i>	<i>M5</i>
3	0.657587963	n/a	0.518214941	0.509086791	33.61982244
5	0.850702545	0.901601007	0.983726984	0.839261745	1.613063299
7	0.949320318	1.12667018	0.976722741	0.906512956	1.636533128
9	0.762272108	0.890131968	0.84667135	0.755234537	0.785741526
11	0.827331113	0.892368403	1.155401395	0.812016885	3.229117797
13	0.959159162	1.058141022	0.967136444	0.915962413	1.729386456
15	0.929235836	1.140204209	1.150282133	0.916996902	1.039279688
17	0.801905371	0.842776903	0.879847027	0.820819272	2.606350771
19	0.957506805	1.102165855	0.959401186	0.917516028	1.70533748
21	1.116970481	1.411693911	1.427118796	1.064142449	1.203876984

Table 39: Average absolute difference between outputs, Trainer/Mediated Output

It is important to note two experiments within this section. The Wang-Mendel controllers functioning with 3 sets experienced problems with two mediation engines. In the tests with a logarithmic function as core of mediation, the agent started normally but then experienced problems that did not allow it to continue. In the case of the controller that was exposed to a mediation of type cube root instead did complete the tests, but was shifting its heading rapidly from one side to the other. If the agent were to operate in a real environment, it could not have been physically possible to complete the tests.

Tables 39 and 40 show the difference between the agent's expected path and the actual path that was taken. If we compare these values to the ones recorded during the preliminary tests, we see that, in the majority of the cases, the agent with Fuzzy Mediation performs better, by showing on average a lower difference. This is not true especially for the Wang-Mendel controller based on 3 sets. This type of controller performs in very similar ways to the preliminary runs. We compare the test of the logarithmic engine with the greedy training, where the "just enough to function" approach to training of the agent was similar to the run of the agent in this section. Since

	<i>M1</i>	<i>M2</i>	<i>M3</i>	<i>M4</i>	<i>M5</i>
3	0.166075739	n/a	0.015558843	0.098498395	0.143920413
5	0.011549945	0.014123278	0.043394838	0.005800359	0.111065528
7	0.08553906	0.080231785	0.041456386	0.07330765	0.113300394
9	0.096699297	0.149351641	0.11243614	0.101637491	0.039836645
11	0.027569965	0.025809044	0.028260324	0.022095325	0.259697977
13	0.115289845	0.102164068	0.017934324	0.087337878	0.035627158
15	0.068979832	0.0513622	0.062385053	0.10790814	0.015440525
17	0.06356347	0.067611014	0.013147392	0.031344037	0.057266212
19	0.098963035	0.111696088	0.047354961	0.057063067	0.021207922
21	0.09845797	0.187985599	0.228074056	0.082817151	0.013798881

Table 40: Standard deviations associated to values in Table 39

the agent was unable to keep going shortly after starting, we can say that it barely had enough rules to fill the important components of the FAM, but the associations were not enough to push the agent through at least one loop. Perhaps, exercising less tight control, the FAM may be exposed to more examples, as the agent is exposed to wider turns, thus more rules to learn.

Finally, we review the information about the average mediation weight, reported in Tables 41 and 42. The data indicates and confirms that the findings in preliminary tests

	<i>M1</i>	<i>M2</i>	<i>M3</i>	<i>M4</i>	<i>M5</i>
3	0.051753052	n/a	-0.164499451	-0.220261745	0.99754481
5	0.359728912	0.080678206	0.524930616	0.42651133	0.461887395
7	0.492531268	0.275630151	0.660261371	0.526306556	-0.159109522
9	0.494606714	0.835798534	0.835865886	0.432146926	0.86220174
11	0.333094972	0.029692105	0.55450172	0.371461626	-0.680352025
13	0.502858064	0.225385935	0.659126335	0.54074085	-0.135325997
15	0.818732258	0.820938795	0.91934512	0.822085056	0.90888542
17	0.333975344	0.033074966	0.528235543	0.38242152	-0.648754562
19	0.489181452	0.249795963	0.659272922	0.542124302	-0.078141946
21	0.879761249	0.824732397	0.888638318	0.83723199	0.896842766

Table 41: Average mediator weight

	<i>M1</i>	<i>M2</i>	<i>M3</i>	<i>M4</i>	<i>M5</i>
3	0.376895729	n/a	0.007609726	0.096119528	0.005489971
5	0.045471403	0.139790982	0.011180285	0.102154983	0.056683789
7	0.062863737	0.042244599	0.066399149	0.078065591	0.095419377
9	0.167249124	0.287288609	0.197705674	0.177798601	0.26155073
11	0.021820949	0.016587061	0.038461627	0.023145623	0.072031872
13	0.082918228	0.070041382	0.051960366	0.107087605	0.109660068
15	0.106275254	0.095060744	0.109853639	0.189207668	0.075075532
17	0.044816938	0.014450254	0.033912239	0.031202935	0.020423724
19	0.096250695	0.083061988	0.067151902	0.082121503	0.086892427
21	0.139221	0.181753428	0.121272868	0.136216096	0.139036219

Table 42: Standard deviations associated to values in Table 41

about the mediation engines hold true in this case as well. We see that, in general, the linear engine behaves in a manner that can be considered middle of the way. The logarithmic and cube root engines show the tightest control in most cases, where the exponential and the cube ones are more relaxed.

Through all these experiments it is important to note that, in the majority of the cases, a 3-set Wang-Mendel controller coupled with an expert controller still manages to learn some associations, but, given the preliminary tests, it performs very well. Such scenario can be placed in the same category of a robot that malfunctions. With Fuzzy Mediation in place, although the AI controller could not pilot the robot correctly, the overall performance of the agent still stays within acceptable limits. This would mean that, a faulty AI application that guides a robot might not destroy it, but just place it in a situation where a safety mechanism needs to be triggered, safely bringing the robot to a stop. This is an example of an application of Fuzzy Mediation either as a metric to evaluate the performance of a novice controller, or as a keep-safe mechanism to reduce

the possibly bad outcomes of faulty AI.

5.5. Fuzzy Mediation with a Simple Flight Simulator

The final step of this research is to study the interaction of people and Fuzzy Mediation in practical applications. For these experiments we will use the simple flight simulator environment described earlier in this work. We will study two distinct domains that may benefit from Fuzzy Mediation. The first is the acquisition of a new motor skill through shared control of a common object. The second is the transfer of motor skill from the user's dominant hand to the non-dominant one. The Institutional Review Board at Towson University reviewed the studies, and the approvals are reported in Appendix A. The questionnaires used in the pre-study as well as post-study phases of the experiments are reported in Appendix A as well. The responses to the questionnaires of the first and the second study are also reported in full in Appendices B and C.

We worked with 21 people, both males and females. The students were of different academic backgrounds, but all enrolled in Computer Science classes. For the first section of the experiments we had 12 experimental runs and 9 control ones. Out of the 9 participants who were part of the control group, we can count only 8, as the data from one of them was corrupted. For the second section instead we had 10 participants, 5 were exposed to Fuzzy Mediation, while 5 were used as controls.

These experiments are aimed at giving us a preliminary insight into the use of Fuzzy Mediation as a learning tool. We reviewed earlier concepts that deal with humans and learning, and Fuzzy Mediation aims to the creation of a novel framework where subjects

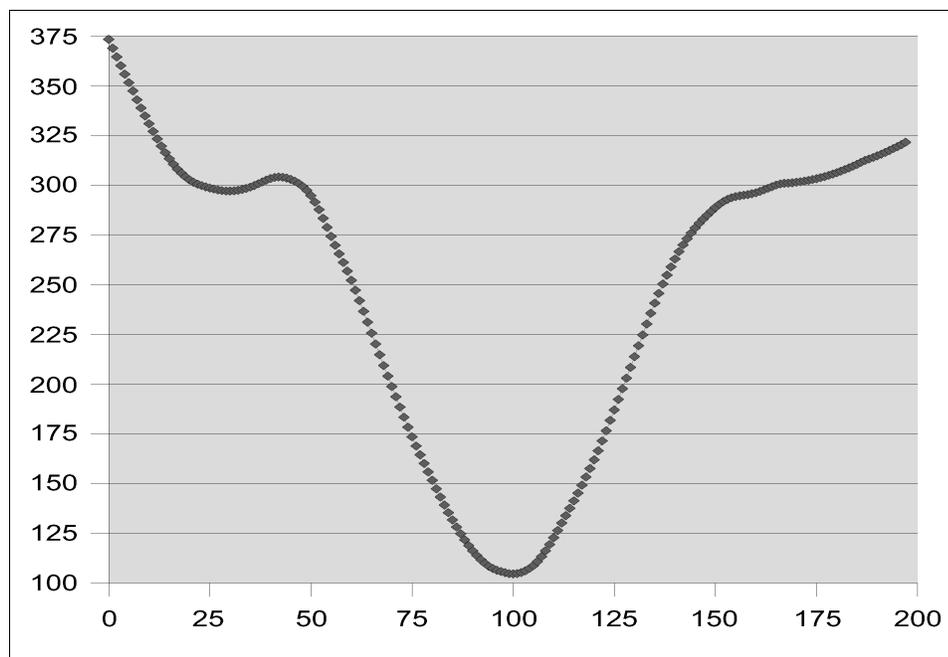


Figure 42: Ideal path, distance from the center of the airfield

can learn in a controlled environment, increasing the exposure to the object to be controlled. The idea of the researcher is to place a person at the controls of a vehicle earlier in the learning process, with the careful supervision and guidance of an expert user. In the case of a simple flight training environment, a student pilot may require a certain number of hours of ground school and simulators before being placed at the commands of a small aircraft (14 C.F.R. Part 61). With this framework, we believe it is possible to shorten those times, thus making training more effective and shorter.

5.5.1. Shared Control

The first set of experiments of Fuzzy Mediation related to the interaction of humans takes place in a simple flight simulator. The object of the game is to control the object within the environment shown in Figure 26 and let it touch certain checkpoints. The checkpoints will turn from red to green. The target of the simulation is to let the object, a

<i>Participant ID</i>	<i>Type</i>	<i>Run 1</i>	<i>Run 2</i>	<i>Run 3</i>
003/H1	Experimental	4	4	4
004/H1	Experimental	2	5	3
005/H1	Experimental	1	4	0
006/H1	Experimental	0	5	0
008/H1	Experimental	2	4	0
010/H1	Experimental	4	5	5
012/H1	Experimental	3	5	2
013/H1	Experimental	2	5	3
014/H1	Experimental	0	5	0
015/H1	Experimental	1	5	3
020/H1	Experimental	3	5	4
021/H1	Experimental	1	5	3
001/H1	Control	0	3	5
002/H1	Control	4	4	4
007/H1	Control	3	5	4
009/H1	Control	2	5	4
011/H1	Control	4	5	5
016/H1	Control	3	5	5
017/H1	Control	n/a	n/a	n/a
018/H1	Control	0	3	1
019/H1	Control	3	5	4

Table 43: Performance of the participants by type

small plane affected by crosswinds, navigate through the checkpoints and land on the designated landing strip.

We recorded the distance of the airplane from the center of the airfield. This measure allowed us to trace a graph of the performance that shows the correct path that any participant should have obtained on a good run, touching all the checkpoints. Such graph is shown in Figure 42.

Table 43 reports the performance of each participant by showing how many checkpoints were touched along the way for each run, keeping in mind that the most

	<i>Run 1</i>	<i>Run 2</i>	<i>Run 3</i>
Average	1.916666667	4.75	2.25
St. Dev.	1.378954369	0.452267017	1.815338686

Table 44: Average performance of the experimental group

	<i>Run 1</i>	<i>Run 2</i>	<i>Run 3</i>
Average	2.375	4.375	4
St. Dev.	1.597989809	0.916125381	1.309307341

Table 45: Average performance of the control group

number of checkpoints a subject should touch within a single run is 5. This table also shows which subject belonged to which group. The performance of subject number 17 cannot be analyzed because the data was compromised during post-study processing. Table 44 shows the average performance of the experimental group, and Table 45 shows the one of the control.

There is some distinction between the experimental and the control groups when it comes to the first run, the one performed by the participants in order to set the baseline, with members of the control group performing better than the subjects later exposed to the shared control. We can also see that the third run of the control group showed a higher rate of success. The one element that jumps to the attention is that, even if a subject performed extremely poorly in the first and third run, the second run, the one with shared control, reported a success of 80-100%. The cause of the airplane touching only 4 checkpoints out of 5 during sessions with shared control is to attribute to the expert controller's decreased attention while trying to help the novice controller learn how to navigate in dual mode.

The pre-study questionnaires, reported in Appendix B of this work, showed a mixed variety of people taking part in these experiments. Very few people had knowledge of the navigation of an object in a moving fluid, given the fact that only two people admitted to knowing how to fly planes or commandeering boats. More people reported that they knew some basics of flight or navigation. Also, most people reported that their eye-hand coordination was quite good, above 7 on a scale of 1 to 10, with 1 being very poor and 10 very accurate. The majority of the people in the study also reported that they play video games 1 to 2 and 3 to 5 times per week. Only two cases reported of playing video games 6 to 10 times per week. Also, two people reported playing no video games at all. When asked if they prefer to read a manual before starting to play or if they just play and figure things out as they go, the overwhelming majority of the participants reported to favor the second option. When asked the type of peripheral that they use the most, the majority of the subjects reported using the keyboard and the mouse. Few reported also using a joystick. Action and strategy games are the ones that most people seemed to favor as genre. The age of the participants fell for the majority of the cases to undergraduate-level college students. A few participants were over the age of 25. The sex of the participants was predominantly male, with only four females taking part to these experiments. It is important to note that one of the participants was trained to pilot submarines, thus we will evaluate carefully his thoughts and opinions of the system.

Now we will analyze the performance of the two groups. It is important to note that the measure used to track the subjects was the distance from the center of the airfield as well as the distance from each of the waypoints. In order to establish a baseline, we can

<i>Waypoint 1</i>	<i>Waypoint 2</i>	<i>Waypoint 3</i>	<i>Waypoint 4</i>	<i>Waypoint 5</i>	<i>Center of Field</i>
334.227238	267.3974526	242.1842894	262.5217624	326.7302414	247.5389158

Table 46: Ideal run

review the performance of an ideal run, reported in Table 46.

The performance of the control group was as expected. On the first run, the subjects showed problems following the path indicated by the waypoints. In all cases the distances from the waypoints as well as from the center of the airfield were quite different from the ones expected. On the third run they showed results that resembled very much the ones in Table 46. Table 47 shows the performance of the control group.

The experimental group showed a similar pattern for the first run. We can then see

	<i>Waypoint 1</i>	<i>Waypoint 2</i>	<i>Waypoint 3</i>	<i>Waypoint 4</i>	<i>Waypoint 5</i>	<i>Center of Field</i>
First run	310.3449238	245.2278532	231.378556	273.3767291	354.6836668	233.3480006
Third run	336.984523	267.9402477	240.0133458	259.6379642	323.9029686	243.7777673

Table 47: Performance of the control group

that the performance of this group on the third run did not yield results similar to the ones in the control group, showing performances that were still far from what was desired.

Table 48 shows the performance of the experimental group.

What is most significant in these experiments is the fact that the second run of the experimental group reported results that were very similar to what was desired. This

	<i>Waypoint 1</i>	<i>Waypoint 2</i>	<i>Waypoint 3</i>	<i>Waypoint 4</i>	<i>Waypoint 5</i>	<i>Center of Field</i>
First run	309.2252481	252.6232062	243.5507797	283.4568769	361.0019658	245.7833579
Third run	327.6511823	257.5521726	237.4406845	268.8189105	342.3589977	240.8024209

Table 48: Performance of the experimental group

<i>Waypoint 1</i>	<i>Waypoint 2</i>	<i>Waypoint 3</i>	<i>Waypoint 4</i>	<i>Waypoint 5</i>	<i>Center of Field</i>
336.2259648	268.9620602	241.4510141	260.6963485	322.8745814	245.0013877

Table 49: Performance of the experimental group during the shared control run

shows that, although the simulator used may not be such an ideal learning tool for motor actions, Fuzzy Mediation allows for a safe, and perhaps ideal, handling of the vehicle being controlled, although two people are at the controls. Table 49 shows the performance of the experimental group in a shared-control setting.

Since Fuzzy Mediation seems to work quite effectively in situations where one of the controllers could not complete the assigned task, we will now review one of those performances in detail. The subject that we will review is number 6. This person was not able to touch any checkpoints, and was unable to complete the simulation during the first run. On the third run, the subject was able to complete the simulation, although no checkpoints were touched. Figure 43 shows the distance from the center of the airfield through the first run.

Figure 44 shows the same measure on the third run. We should note that the graph shown does not resemble the one in Figure 42 shown earlier, thus indicating that the simulated plane was all over the airfield, rather than following the expected path.

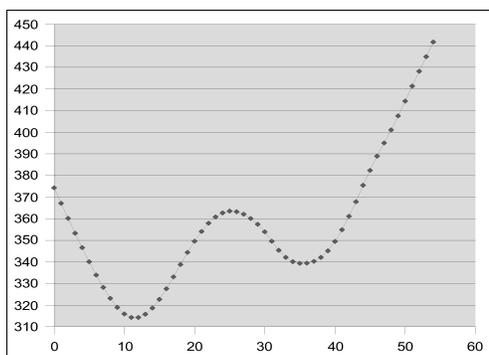


Figure 43: First run of test subject 6

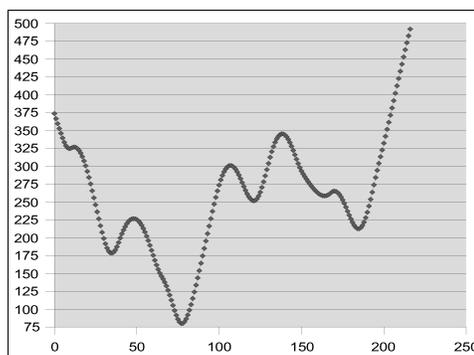


Figure 44: Third run of test subject 6

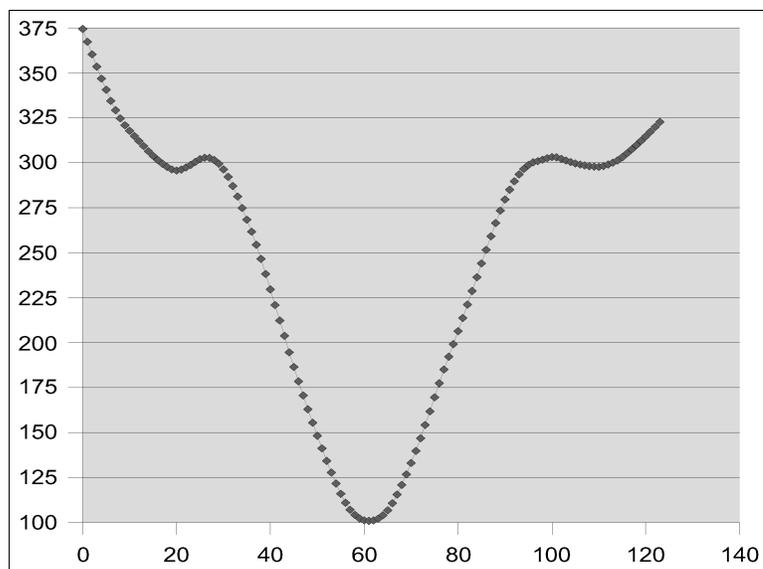


Figure 45: Second run of test subject 6

We should now focus on the second run, the one where two subjects controlled the plane. Figure 45 shows the average distance from the center of the airfield. In this case, the graph reported looks very much alike the desired one, shown in Figure 42.

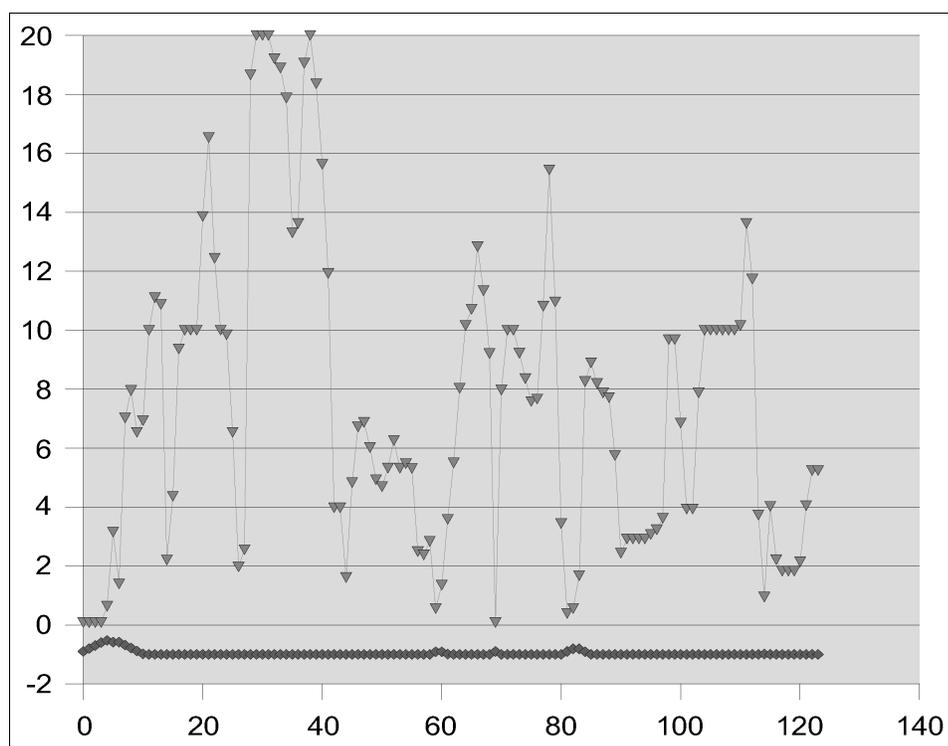


Figure 46: Mediation weight $([-1, 1])$ and abs. difference between inputs (run 2)

Next, we should review how the weight of control shifted through this simulation. Figure 46 shows how control was mediated between the two pilots. We can see that the expert controller was handling the plane most of the time, as the novice was trying to maneuver, without much success.

In the second run, the plane was able to touch all five of the checkpoints, thus completing the simulation successfully, while the single novice controller was unable to do so in both the first and the third runs.

As we reviewed the post-study questionnaires, reported in Appendix B of this work, we saw many reactions among the people in the experimental group. When asked how they felt on their first attempt to navigate the object through the simulation for the first time, most people responded that they were feeling anxious, confused and overwhelmed by the speed or the sensitivity. During the second attempt, people who belonged to the control group reported that they were feeling more relaxed and in control of the situation. Some people who belonged to the experimental group reported to feeling as if they were not in control, but most reported on the speed of the simulation, as it was slower. One participant reported that it was “Easy – trusted the ‘co-pilot’ and myself. Just focused on navigating. Pretended copilot was not there.” Another participant reported that it was “Easier [because] it was slower, but a little intimidating doing a co-pilot.” When asked if they thought they learned something during the second round of simulations, most people in the control group stated that they did feel an improvement, while we received mixed reactions from the experimental group, with many people reporting a little improvement, but not much. During the last attempt most participants who belonged to the control

group reported to be more confident and in control of the object. Participants assigned to the experimental group seemed also confident. Both groups reported about the higher speed and how it threw them off a little during the third attempt. When asked to rate how much they learned from the simulations, both groups reported an average improvement. No group stood out more than the other. Finally, when the participants had a chance to add their own comments, few people responded. The comment that attracted our attention the most was the one that stated: “Might be too simple to really observe or experience improvement.”

Much research is needed to create interfaces that truly convey enough information for one person to learn how to operate a vehicle through a simulator. Especially when the control is shared, there is a great need for transparency and ease of information architecture when addressing what needs to be shared and what not. The creation of successful training programs is highly dependent on the nature of the situation as well as how much information really needs to flow from one controller to another. We do feel that the interface we created was not perceived too well by most users, but it still gave us an insight on how much shared control through means of Fuzzy Mediation can help people completely unable to guide an object through a relatively simple path in safety. Some of the examples shown earlier report that the participants were completely unable to reach the goal in some situations, and they had an even more difficult time in touching the checkpoints. When placed in a situation where control was shared, the performance improved greatly, resembling very much the expected path. The same Fuzzy Mediation framework allowed other participants to successfully perform through the simulation and

gain control of the object.

In conclusion to this experiment, we can say that more work is needed to create an appropriate interface to accommodate for Fuzzy Mediation as a mean of teaching motor control from one person to another, but we can also state that an object can be successfully controlled by two people, in a situation where one of the two could not do it alone.

5.5.2. Bi-manual Control

The second experiment that involved the interaction of people with the Fuzzy Mediation framework involved the analysis of motor skill transfer and analysis in an environment that utilizes bi-manual control as a learning tool. Since previous tests with human controllers showed a lack of a proper interface, we felt that bi-manual coordination was the next step to try.

In bi-manual coordination the majority of the feedback as far as motor control alignment between a trainer and a trainee happens at the cognitive system level within the person. We do not need elaborate interfaces to display that the trainee is under-steering or over-steering as the agent needs to change direction.

For this experiment we used the same environment as in the previous tests. Since the environment stays unchanged we can assume that the desired path will resemble the one shown in Figure 42 above.

<i>Part. ID</i>	<i>Type</i>	<i>Run 1</i>	<i>Run 2</i>	<i>Run 3</i>	<i>Run 4</i>
001/H2	Experimental	5	5	4	5
002/H2	Experimental	4	4	5	5
003/H2	Experimental	5	4	5	5
004/H2	Experimental	4	4	3	5
008/H2	Experimental	5	5	5	5
005/H2	Control	5	4	5	5
006/H2	Control	5	5	5	5
007/H2	Control	3	5	5	3
009/H2	Control	5	5	5	5
010/H2	Control	4	5	5	5

Table 50: Performance of the participants by type

Table 50 reports the type of the participants along with the number of checkpoints that were successfully touched. The first run was performed by the subject using their dominant hand. The second run instead was performed using the non-dominant hand. The first two runs were used as baselines to understand the performance of both hands and use it as reference later on in the analysis. The third run was different for each experimental group, the control group was required to perform another test with the non-dominant hand alone, while the experimental grouped was allowed to perform the test with both hands, the dominant one in the role of the trainer, while the non-dominant one in the role of the trainee. The fourth run was performed with the non-dominant hand to assess any improvements compared to the second run. Also in this case there is no clear advantage to using Fuzzy Mediation by just reviewing the number of checkpoints touched along the way. As we will see later though, the actual path of the airplane and the mediation weight show the significant role that Fuzzy Mediation plays in these experiments.

	<i>Run 1</i>	<i>Run 2</i>	<i>Run 3</i>	<i>Run 4</i>
Average	4.6	4.4	4.4	5
St. Dev.	0.547722558	0.547722558	0.894427191	0

Table 51: Average performance of the experimental group

Many subjects who took part in the first experiment were also part of this test. As the participants are a subgroup of the first set, the results from the questionnaire administered before the simulations are very similar to the ones reported above. It is important to note that also in this case one of the participants was trained to pilot submarines, thus we will evaluate carefully his thoughts and opinions of the system.

When we compare the performance of the two groups, we can see little difference. Given the small amount of participants, we can only use these data as an indication of a possible trend, rather than a certainty of the users preferring one method to the other. When we review the performance of the experimental group by looking at the average number of waypoints touched per run, reported in Table 51, we can see that, overall, the performance on the fourth run, the third using the non-dominant hand, surpasses the performance of the second run. Moreover, we can see that the average number of checkpoints touched is higher than even the control run, the one performed with the dominant hand.

As we compare these results with the ones obtained by the control group, reported in Table 52, we can see that the performance of this second group is not as promising as the

	<i>Run 1</i>	<i>Run 2</i>	<i>Run 3</i>	<i>Run 4</i>
Average	4.4	4.8	5	4.6
St. Dev.	0.894427191	0.447213595	0	0.894427191

Table 52: Average performance of the control group

	<i>Waypoint 1</i>	<i>Waypoint 2</i>	<i>Waypoint 3</i>	<i>Waypoint 4</i>	<i>Waypoint 5</i>	<i>Center of Field</i>
Run 1	334.227238	267.3974526	242.1842894	262.5217624	326.7302414	247.5389158
Run 2	354.6109407	284.103733	248.5233873	257.9911693	309.5198885	251.3878119
Run 3	347.3526004	278.0711797	249.9106659	265.0107441	319.6145595	253.1322704
Run 4	341.8990694	273.4879071	243.4023336	259.225989	318.9165862	247.738979

Table 53: Ideal run

one in the first. We can see that the second trial using the non-dominant hand, run 3, records an improvement over the first, but the third, run 4, shows a decline.

When compared to the base run, the one performed with the dominant hand, both groups seem to have performed better by the last run. This is probably to attribute to the fact that, by performing a new motor task with their non-dominant hand, they are forced to pay more attention to the movements, rather than some sort of autopilot mode that engages as the confidence in their dominant hand is tested.

Also in this case we should keep an ideal run as a reference for the next analyses. Table 53 shows the performance of one subject that seemed the best. Also in this case, we review the average distance between the position of the airplane and the five waypoints as well as the average distance from the center of the airfield.

First, we will review the performance of the control group, reported in Table 54. We can see that the performance of this first group was not too close to the one of the ideal run. We can see that the values fluctuate, but there is no distinct improvement as the non-

	<i>Waypoint 1</i>	<i>Waypoint 2</i>	<i>Waypoint 3</i>	<i>Waypoint 4</i>	<i>Waypoint 5</i>	<i>Center of Field</i>
Run 1	336.3866478	268.3508808	241.2784449	261.0594869	323.8061177	245.6811206
Run 2	335.2999865	267.2775895	238.9770011	258.671905	322.2560889	243.2963102
Run 4	330.3451326	263.2301295	238.8474837	261.2298763	326.9628061	242.9967489

Table 54: Performance of the control subjects

	<i>Waypoint 1</i>	<i>Waypoint 2</i>	<i>Waypoint 3</i>	<i>Waypoint 4</i>	<i>Waypoint 5</i>	<i>Center of Field</i>
Run 1	337.9874949	267.7914758	237.8384416	254.1918495	319.8878093	242.8946584
Run 2	326.1371635	261.896416	238.2616811	263.6240728	335.3498809	244.0084503
Run 4	343.7097964	272.1120321	241.6174969	257.7103766	318.2717449	247.2802302

Table 55: Performance of the experimental group

dominant hand is exposed to more trials. The third run of the control group was performed using the non-dominant hand and was not reported, as it was just a test run to get the user more familiar with the controls.

The experimental group instead shows a definite improvement, as is reported in Table 55. The values of this group from the first run performed using the non-dominant hand, run 2, resemble the ones recorded for the control group. But when we analyze run 4, the one monitored for improvement and performance, we can see that the values are close to the ones reported in the ideal run.

We should also review the values recorded during the third run that involved the subject's non-dominant hand, reported in Table 56. In this particular run the subjects were allowed to use both hands, with the role of the expert assigned to the dominant hand, and the novice being the non-dominant one.

In this table we can notice that the performance was not extremely close to either the subject's base run, performed with the dominant hand, or to the ideal run. It is important to note though that, in some the cases, this run saw an improvement in the number of waypoints touched, when compared to the first run performed using the non-dominant hand.

	<i>Waypoint 1</i>	<i>Waypoint 2</i>	<i>Waypoint 3</i>	<i>Waypoint 4</i>	<i>Waypoint 5</i>	<i>Center of Field</i>
Run 3	320.5948795	257.1776004	243.4967386	273.0755147	344.9251335	248.4415463

Table 56: Experimental group, run 3 (bi-manual control)

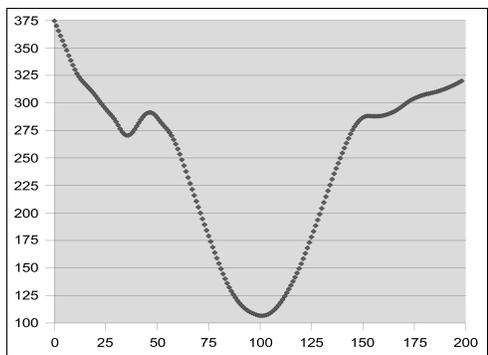


Figure 47: Performance of subject 4 in run 1

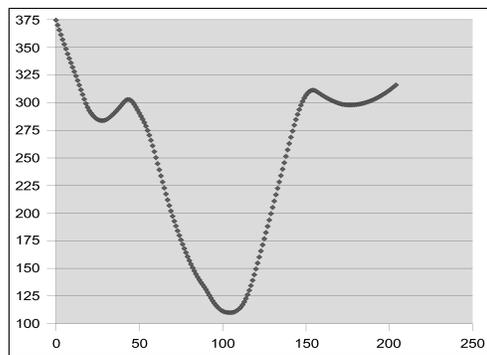


Figure 48: Performance of subject 4 in run 2

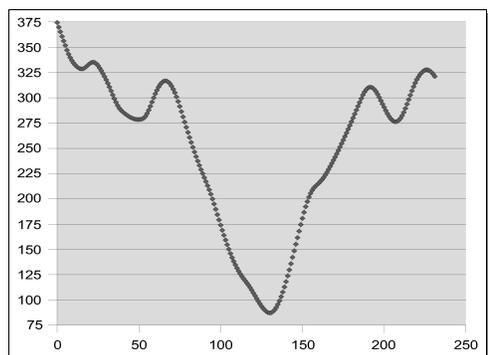


Figure 49: Performance of subject 4 in run 3

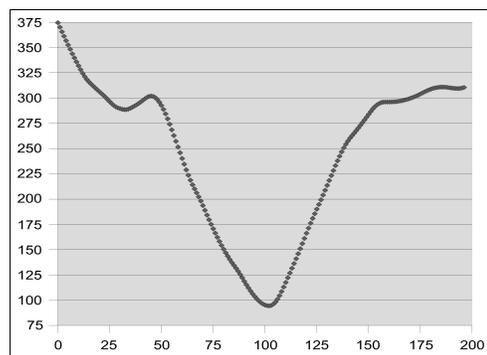


Figure 50: Performance of subject 4 in run 4

These results seem to suggest that, in the case of bi-manual coordination, Fuzzy Mediation may have an effect on the transfer of motor skills higher than when transferring the same knowledge from one person to another. A possible explanation of this trend is that the internal muscle-skeletal feedback systems play a role in letting the brain know exactly what kinds of motor actions the limbs are performing (Martini, 2005). This type of internal comparison and regulation system makes any type of interface not as necessary as in the case of shared control between two subjects, where only visual feedback can help the novice learn from the expert. It is important to remember though that the data reported in this section of the work should be used as an indicator of a possible trend, as not enough subjects were used to make these experiments statistically

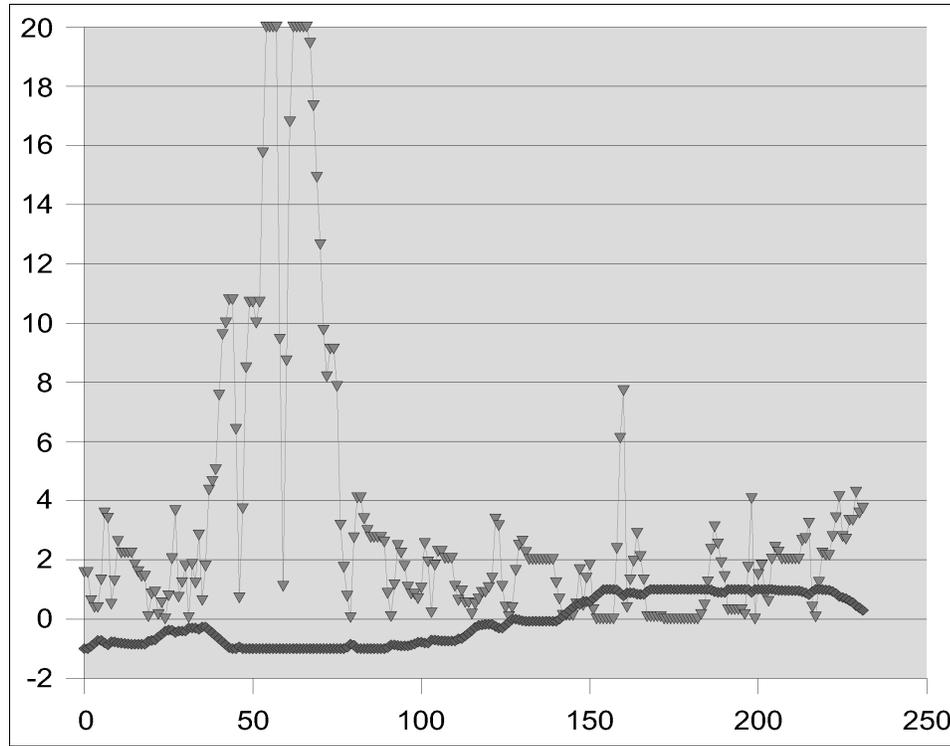


Figure 51: Mediation weight $([-1, 1])$ and abs. error between controllers (run 3)

significant.

Also in this case, we should analyze the performance of a single subject and evaluate the Fuzzy Mediation framework from a slightly different point of view. As stated earlier in this work, one possible outcome from all the experiments performed is the introduction of Fuzzy Mediation as a metric to evaluate closeness in performance. The case of bi-manual control is perhaps the most interesting, as the biological feedback mechanisms built into the human body make for a perfect coordination equipment.

Figures 47 through 50 show the performance of the subject for the four runs. We can see that all resemble more or less the graph shown in Figure 42, thus the performance of this subject was always satisfactory.

What is most important in this case is Figure 51. Here we can see that, at the

beginning of the simulation, the subject was not coordinated, as the weight of control was leaning towards the negative numbers, rather than the positive ones. Towards the middle of the run though the subject showed a great improvement, allocating most of the control from that point on to the non-dominant hand, as it is reported by the weight of control being in the positive region.

In the post-study questionnaire, many participants felt awkward about using their left hand in the simulation, especially in the experimental group. During the run performed by the experimental group with dual control, there were mixed reactions, some felt “Good about a somewhat steady green indicator”, referring to the gauge that would turn brighter green as the trainee was more in control, and whiter as the trainer controlled the object. Also, one participant stated that it was “Frustrating. Focus was right hand.” And another participant wrote the following: “It seemed difficult at first to coordinate both hands but it was easier towards the end.” The control group reported feelings that are consistent with a typical training run, where they felt more in control. Most participants felt that they learned from the training run, one person though wrote that there was no improvement, as the performance of the earlier run was already fine. The fourth run, the third using the left hand, reported all participants much more in control and satisfied with their performances. On average, the control group and the experimental one reported similar values as far as how much was learned during these experiments. The additional comments were just two, and all belonging to the control group. They did not state anything relevant to the framework or the training mechanisms.

Also in these sets of experiments we noticed that there was no clear improvement in

the transfer of motor skills from the dominant hand to the non-dominant one, Fuzzy Mediation has great potential as a training tool for bi-manual tasks. It also shows that it is a powerful metric that displays how close the two hands are performing. This type of metric is extremely important in fields where bi-manual coordination is the key element to the success of tasks.

5.5.3. Observations

The poor performance of Fuzzy Mediation as a training tool can be probably found in the fact that the simulators used in these experiments were not performing tasks that were extremely complex, and the interfaces did not convey as much information as needed. These findings are in line with the research of O'Malley et al. (2006), showing that Fuzzy Mediation can easily be used as a metric and an evaluation of performance, but the interfaces need more work to reach the level of applications used for training purposes.

It is very important to note that a factor that may have played a role in the low clarity of the results may be the short time the subjects were given to familiarize with the training environment. Bennett and Pilkington (2001) stress how important it is for a subject to feel comfortable and get used to the settings through which learning occurs. This concept requires careful consideration, in order not to interfere with concepts of habituation discussed by Kandel, Schwartz and Jessell (2000).

5.6. Interview with a Pilot of Submarines

One subject who took part to both experiments that studied human subjects

Question 1: Does the protocol for teaching how to pilot submarines involve simulators? If so, how long is the training through a simulator?

Yes, there is an entire building utilized for the purpose of teaching and training members of the submarine force how to pilot a submarine. [It is] approximately 6 weeks long for the basic submarine school. During that time members will learn the basics of the submarine and its components. Everyone then will participate in several piloting session[s] with a simulated control panel.

Question 2: Does the protocol for teaching how to pilot submarines involve practices of bi-manual control? If so, what kinds of tasks do they train your non-dominant hand to perform? How about your dominant hand?

The protocol for teaching submariners how to pilot submarines involves the usage of both hands at the same time and never one hand at any given time. This is even the case when someone is about to relieve you of your time piloting. In this case, the person relieving will place one hand on the steering mechanisms and the relieved person will keep one hand on until the reliever has full control.

Question 3: Do you think the training program for submarines would benefit from shared control environments? If so, how?

Yes, it will help one learn when and where to apply a certain amount of pressure and help them to understand how to guide the submarine on reaction to the other [person's] skill.

Question 4: Do you think the training program for submarines would benefit from environments where your non-dominant hand is trained by your dominant hand? If so, how?

Yes, it will help train their non-dominant hand and will ultimately result in better steering and handling of the submarine in rough waters.

Table 57: Interview with a pilot of submarines

interacting with the Fuzzy Mediation framework was a pilot of submarines. In the interest of finding out more about first-hand experience with real-life shared and bi-manual control systems, we created a short questionnaire to assess the need for a novel environment and mediation structure that would help train submarine pilots. The questionnaire is reported in its entirety in Table 57.

As the answers of this subject reveal, shared as well as bi-manual control systems are very important in the field of submarines. Also, much time is spent in training for such systems. We believe that Fuzzy Mediation may help either shorten such times by playing a pivotal role in assessing the performance of the pilots, and perhaps power some of the simulators that deal with shared control systems as well as check-rides.

5.7. Summary

In this chapter we reviewed many experiments that illustrate the true power of Fuzzy Mediation as an innovative and dynamic extension to traditional Information Fusion techniques. First we reviewed the interaction of static controllers trying to control an agent. Then we looked at the interaction between a static expert controller and an A.I.-based pre-trained novice. The third set of experiments analyzed the engine that controls the mediation, allowing the Fuzzy Mediation framework to behave more or less conservatively, maintaining the majority of control either to the expert or to the novice, depending on the setting in which this algorithm is used. Next, we reviewed experiments that had an expert controller interacting with a blank A.I.-based controller, situation in which the novice was bound to learning from the expert in order to control the agent. Finally, we reviewed the performance of humans using the Fuzzy Mediation engine in order to either allow two different users to interact with the agent or to let the same user control the agent, but with the dominant hand acting as expert and the non-dominant one as novice.

In the next chapter we will analyze once again the high-level view of Fuzzy Mediation, now that we have a better idea of the technical operations and the correct fusion of information that the program generates.

6. Discussion

Nowadays many domains thrive on the interaction between humans and machines. Fuzzy Mediation seems the next logical step in order to destroy the barrier imposed by an automated controller that overlooks every action taken by a human before it is actually put in action. These types of systems can be found especially in the new fly-by-wire airplanes developed in recent years.

Not only will this algorithm create a middle ground for the control of vehicles, but it will also create a framework where learning is not executed by passive observation of someone else's performance, and at the same time is not as drastic as putting a novice pilot in control of a plane for the first time. Fuzzy Mediation naturally uses positive reinforcement as the key that will enable trainers and trainees to interact through the first phases of the learning period, as to give the student a sense of control, but at the same time, not to relinquish control completely.

It is important to note that Fuzzy Mediation should allow for a paradigm shift in the type of learning environment when both operators reach a level of familiarity with the controls that is equal. At that point the Fuzzy Mediation algorithm will not use positive reinforcement as leverage to learning and operating a vehicle, but it will use collaborative learning, where two agents need to interact successfully to accomplish a task.

Other fields of application include monitoring systems that learn from the operator of a vehicle. The vehicle will observe the operator for a certain lapse of time, and it will continue doing so for the majority of the operational time of the machine. The vehicle will maintain the role of a trainee as long as the vehicle is operated within certain guidelines. When the boundaries for safe operations will be broken, then the trainee will

be able to replace the trainer in operating the vehicle, effectively switching roles. The roles will be switched back to normal when the trainer, now trainee, will show that s/he can operate within safety guidelines for a certain amount of time.

We should also note that it is important to study the dynamics that form and evolve when multiple controllers are interacting through Fuzzy Mediation. We will continue to explore the interaction between agents and the emergence of new behaviors in homogenous and heterogeneous multiagent systems, as described in Arai & Ishida (2004) and in Trajkovski (2007).

6.1. Possible Applications

Fuzzy Mediation seems to be a very promising approach to a new architecture for shared control environments. Given its nature, its applications can exist virtually in any niche where there is the possibility of having two separate controllers interacting to guide a single vehicle. In this section we will discuss some applications that seem ideal for this algorithm. Some of the material discussed here is not necessarily easy to implement at this point in time, as some of the technologies may be missing. The next topics are aimed at building a roadmap for Fuzzy Mediation and its possible future applications.

6.1.1. Automotive technologies

The automotive industry is taking giant leaps each year towards more and more automated cars. Some are still at the development stage, such as cars that follows the lanes on the road (Tsugawa et al., 2001; Yang et al., 1999), others are already in production, such as the latest Lexus that parks itself by using sensors and cameras

(Lexus, 2007). Also great attention was drawn to the automation of the process of driving by the races in the desert that are held yearly, where only automated driving systems are allowed to let the cars involved race from start to finish (Niccolai, 2004).

It is not too far of a stretch to envision cars in a near future that will drive without the need of a human controller, as written by Norman (2005). It is at this point though that the concept of shared control becomes extremely important, and Fuzzy Mediation is a viable framework to create such environment.

Also, people becoming sick as they drive generate many accidents (Ysander, 1970). It is important to realize that there will soon be technologies that can analyze a driver's performance and compare it to a set of notions about that very driver's driving style learned from the car as it operates. Fuzzy Mediation may play an important role in the case when the driver should become impaired from an illness or from other reasons, and the driving behavior change quite a bit from the normal pattern. At that point the car could take over and safely come to a stop, rather than become a weapon and risk people's lives as the driver loses control.

This research will be extensible to other fields, such as the creation of an autopilot system that uses "drives" as parameters to control an automated robot, as explained in Trajkovski et al. (2005). Such robot will also be able to detect and avoid obstacles in an unfriendly environment using the structure described in Trajkovski, Stojanov & Vincenti (2005).

Some applications of this algorithm are not as obvious as it may seem. An environment where control is shared may be as obvious as a dual-control airliner. Such environment may also be more subtle, where operations are constantly under the scrutiny

of an automated controller, but a human performs the majority of the operations, such as the controller of an anti-lock braking system in a car. Fuzzy Mediation would thrive in this environment as it would mediate between the force exerted by the human controller, under the guidance of the computer that ensures the brakes will not lock, thus putting the vehicle in a dangerous situation.

6.1.2. Aeronautical Technologies

Modern airplanes are working on a fly-by-wire type of technology. Little room is left for hydraulics, and especially in the case of Airbus, full automation is in place (ISA, 2006). This excess of automation leads to many advantages, such as lighter airplanes, but also to many drawbacks, such as the on-board computer being in charge of approving or denying any possible input given by the pilots. Also, in this case Fuzzy Mediation may play a pivotal role in allowing shared control.

Different airplane manufacturers have different visions of how a flight deck should be structured, as well as the kinds of controls that the pilots should use. Airbus was the first to create a fly-by-wire aircraft that also had dual flight controls independent of each other. This means that, when the pilot in command performs a left turn, the stick that belongs to the other pilot does not turn. Other airplane manufacturers, although still use the fly-by-wire technology, preferred to have linked control systems, where the motions of one pilot are reflected in the controlling stick of the other. The policy that Airbus adopted seems perfect for the framework created by Fuzzy Mediation, where two pilots can interact at the same time with the same aircraft, while keeping the operations completely safe. At this time, only one pilot can interact with the flight controls. In some

cases the action of the disabled controller may actually make the problem worse as the airplane was not designed to take in two inputs (Fiaschetti, 2007).

Another component of the aeronautical industry is that all pilots must go through a certification process where FAA officials evaluate the flight performance. These events are called “Check-rides” (14 C.F.R. Pt. 61). Another possible application of Fuzzy Mediation could be the automation of check-rides by allowing the pilots to fly through pre-defined patterns with the supervision of this algorithm. The constant comparison between expert and novice behaviors would allow for the aircraft itself to monitor the pilot’s performance and compare it to the FAA-approved expected path. Also, this mechanism could be integrated on all planes, not as a check-ride evaluation system, but as a mechanism that evaluates the airplane’s performance and flight path, communicating constantly to a central information system that may help prevent hijackings, or at least become aware of them as soon as the first signs of deviation from the expected flight path are recorded.

6.1.3. Medical

Fuzzy Mediation has a great potential also when applied to the field of medicine. Nowadays much effort goes into research for telemedicine solutions, so that a doctor does not need to be present in the operating room in order to perform a surgery on a patient (Karamanoukian et al., 2003). The algorithm proposed in this work would fit very well in training as well as operations environment in a telemedicine framework. Also, research recently published shows that surgeons perform better when they play video games (Rosser et al, 2007). It is only reasonable to assume that, if certain video games were

designed keeping the improvement of motor skills in mind, the dexterity of the players would improve even more. Also in this case, Fuzzy Mediation could play a pivotal role in such games.

Also, rehabilitation plays an important role in the field of medicine. Many people suffer each year from strokes or unilateral paralyses, thus losing control of one side of their bodies. The process that these patients have to undergo from the point of the paralysis on is extremely long, if possible at all, so gain even a minimal functionality back on the side affected. Recent works reported that stroke patients might actually benefit from playing video games (You et al., 2005). Also in this case it is not unlikely that carefully designed video games may also improve the process of recovery. Research in this area from the author of this work has already started.

6.1.4. Training

Joystick-like devices control many types of machinery. One category of machines that has moved its full line of equipment to joystick-only is Caterpillar's work machinery. Excavators, cranes and such are all equipped with operate-by-wire technologies that are controlled by joysticks. The training environment created by Caterpillar focuses heavily on the eye-hand coordination (Caterpillar, 2007). Sometimes bi-manual types of controls operate the machine, other times each hand is assigned a different task, but both need to be performed at the same time. Also in this case Fuzzy Mediation is not only a possible framework that may govern training programs, but also it may control operations of the heavy machinery.

As discussed earlier, some training programs require a certain number of schooling

and simulations before the student is placed at the controls of an object, such as in the training of airplane pilots (14 C.F.R. Part 61). The framework introduced in this work may help shorten the pre-exposure requirements, thus making training more effective and at the same time shortening the program.

6.1.5. Metric

The last possible application that Fuzzy Mediation may assume that we will discuss in this work is its role as a metric. Many applications that we already introduced would use Fuzzy Mediation as metrics to study and analyze how closely two controllers or two hands were performing in a joystick-driven control system. It seems only obvious to use Fuzzy Mediation as a performance evaluation in the check-ride situation described earlier, as well as the Caterpillar training programs. It is perhaps a bit less obvious, but perhaps even more important, to utilize Fuzzy Mediation as a metric in the recovery process of patients who suffered a unilateral paralysis. The creation of a system that would record how closely two hands are performing motor actions through the recovery process may improve the life of a patient, who may perform such exercises at home instead of having to be driven to a rehabilitation center, but also it would allow the physical therapists to easily monitor progress from a central repository of patient data. Work in this direction has already started.

6.2. Future work

As the algorithm presented in this work is very promising, it is also very young. There are many aspects that were brought to light through these experiments that need work and further exploration. This section highlights some of these components.

The first area of improvement that this framework needs is in the creation of better interfaces for learning settings. Much research was done in the field of feedback mechanisms applied to shared control with particular focus on haptics. Steele and Gillespie (2004) performed research in aiding the guidance of land vehicles. O'Malley et al. (2006) and Li, Patoglu and O'Malley (2006) instead looked at haptics applied to learning new tasks. More research is easily available praising the benefits of haptics applied to control in general and especially to shared control (Goncharenko et al., 2004; Griffin, 2003; Griffin, Provancher & Cutkosky, 2005; Griffiths & Gillespie, 2005; Griffiths and Gillespie, 2004) leading us to the conclusion that the algorithm described in this work will benefit as a learning tool by the addition of a feedback based on motor response felt through the controlling medium.

It is also important to assess the validity of Fuzzy Mediation as a mean to learning motor tasks. Such assessment will be easily conducted through the analysis of the closeness of the inputs of the two controllers through a temporal data mining tool. Such tool, reported by Vincenti, Hammell & Trajkovski (2005) and an extension of a conventional Market Basket analysis (Agrawal, Imielinski & Swami, 1993), will highlight changes in the behavior of the novice controller as the subject performs.

The topics discussed in the previous section highlight the potential of Fuzzy Mediation as a metric for performance evaluation. Given the many possible domains to which this process may be applied, it is important to evaluate Fuzzy Mediation carefully in a controlled setting before it is used in live applications, where its malfunction or inadequate setup may lead to faulty operations.

We also discussed the application of Fuzzy Mediation to the automotive and aerospace industry. Given the nature of these vehicles, it is important to test and evaluate the performance of Fuzzy Mediation using either small-scale vehicles or automated drones in order to understand the true implications of shared control applied to these types of machines.

Finally, there is a great need to study Fuzzy Mediation as a collaborative control environment. The experiments reported in this work assume that the level of knowledge of one controller is superior to the one of the second controller. It is inevitable that, as operations continue, the level of knowledge of the two controller will grow to adapt to each other, thus creating a situation where there is no obvious superiority of one person's performance over the other. In this case the teacher/student approach to Fuzzy Mediation will not suffice. There needs to be an extension to this framework so that the two inputs are evaluated equally and the comparison metric is altered to reflect the new paradigm. The author of this work has already started research in this direction.

7. Conclusions

Fuzzy Mediation is a novel ideal created by the author of this work. It was originally drafted by recognizing the need for more effective methods of on-line learning to let humans and artificial agents learn how to operate vehicles in real-life situations. This approach was also created keeping in mind the disciplines of shared control as well as information fusion, in order to truly allow the controllers be placed at the commands of such vehicle, and still retain the safety that is necessary when learning a new motor task. Moreover, part of the original set of goals was the extension of traditional information fusion algorithms to accommodate for a dynamic fusion mechanism that would not rely on pre-defined weights or algorithms that would fix the creation of the single output. And lastly, the creation of a new and innovative framework also demands the creation of a new metric, which can also become a part of the framework itself.

The chapters presented in this work illustrate all of the concepts the authors had in mind at the beginning of this research. We learned about the high-level as well as the technical aspects of Fuzzy Mediation. We also reviewed some simple experiments coupled with some more intricate ones, which saw simulators and agents that displayed dual controllers. Such experiments showed that not only Fuzzy Mediation is an approach that works, but they also reported results that are of indisputable evidence that this algorithm is a key component to on-line learning to systems that may thrive on truly shared controls for either training or operational uses.

Fuzzy Mediation showed the ability to place at the controls of an agent an A.I.-based

system with no knowledge of the environment or the control strategies to be used. Then it placed it in a condition of safety for the agent to perform its line-following duties, still allowing the novice to influence the overall direction when appropriate, and keeping it away from controls and the simulated robot safe when the novice's directions were deemed too different from the expert's guidance. When dealing with A.I.-based agents, we are usually facing a training phase followed by an operations phase. With Fuzzy Mediation agents can skip the formal training phase to blend it into the operations phase, where the agent basically learns and improves its knowledge in full safety.

Although the experiments with human controllers were not as clear as we wished, they were still of vital importance to show the safety mechanism that Fuzzy Mediation assumes when one expert controller and a novice one interacted with the simulation. In one of the cases recorded, the novice could not complete the simulation alone, but in conjunction with the expert, the simulated plane was able to not only complete the desired path, but also pass through the majority of the checkpoints.

During the second set of experiments involving human controllers, we were able to see how important of a metric Fuzzy Mediation can become for shared control environments. In our case, we compared the operations of a human controller faced with the task of piloting the simulated plane bi-manually. We were quickly able to see how the hands were working in different manners, when executing the motor actions linked to controlling the plane. As the simulation progressed, we were also able to see how control shifted from the "expert", the dominant hand, to the "novice", the non-dominant one, thus

showing that the hands were coordinating in motor actions.

We then introduced some possible practical applications of Fuzzy Mediation to real-world situations. Although there are many, the author of this work will continue working in the direction of improving this algorithm and any corollary applications that may be required in order to apply it fully to bioinformatics and medicine.

In conclusion, this work represents the first step towards a new and innovative branch of computing sciences. The many possibilities of further research, improvement and applications make it a very important step towards true shared control, dynamic information fusion, start a new framework for on-line learning and place it as an important metric in dual-controller environments.

Appendices

Appendix A – Institutional Review Board Documents

Informed Consent Form

Dear participant:

My name is Giovanni Vincenti and I am a doctoral candidate for the degree of Doctor of Science in Applied Information Technology at Towson University. The title of my study is “Fuzzy mediation as an improved method towards motor learning”. The concept of fuzzy mediation involves the collaboration of two people trying to manipulate an object with control dynamically shifting from one controller to the other as the performances of the two become more similar.

If you decide to participate in this experiment, you will be requested to follow a pattern in a simulation created by me. In this simulator you will be required to maneuver a vehicle, such as a boat or a small airplane, through a course. I will record your reactions as well as the path that the object you are controlling follows. You will also be required to fill out a questionnaire before you use the simulator as well as one after the study. The data collected through the simulation will be used to evaluate a novel method of cooperative learning environment applied to the acquisition of motor skills.

Participation in this study is voluntary. All information will remain strictly confidential. Although the descriptions and findings may be published, at no time will your name be used. You are at liberty to withdraw your consent to the experiment and discontinue participation at any time without prejudice. If you have any questions after today, please feel free to call me at 410-704-4770, contact me at gvincel@towson.edu, or contact Dr. Patricia Alt, Chairperson of the Institutional Review Board for the Protection of Human Participants at Towson University at (410) 704-2236.

Consent of the participant:

I, _____, affirm that I have read and understood the above statement and have had all of my questions answered.

Signature: _____

THIS PROJECT HAS BEEN REVIEWED BY THE INSTITUTIONAL REVIEW BOARD FOR THE PROTECTION OF HUMAN PARTICIPANTS AT TOWSON UNIVERSITY

Pre-study Questionnaire

Do you know how to fly airplanes or commandeer boats?

Yes No

If yes, explain what kinds of airplanes/boats you are able to commandeer.

Do you know the basics between flying airplanes or commandeering boats?

Yes No

If yes, indicate the extent of your knowledge on a scale of 1 to 10, 1 being that you were exposed to a short introduction to the subject and 10 where you took extensive courses on the subject.

Also, indicate what kind of exposure you have received.

Value: _____

Type of exposure:

Rate your eye-hand coordination on a scale from 1 to 10, with 1 being very poor and 10 being very accurate.

Value: _____

How often do you play videogames?

- 1-2 times a week
- 3-5 times a week
- 6-10 times a week

Date: _____

Participant number: _____

When faced with playing a videogame or learning a new motor task, what do you like to do first?

- Read the instructions manual Start playing and figure things out as you go

What kinds of interfaces do you use to play videogames? (Check all that apply)

- Keyboard
 Mouse
 Joystick

What kinds of videogames do you play with most of the times?

- Action (ex: Quake, Half Life)
 Strategy (ex: Civilization, Rome)
 Precision shooting (ex: Call of Duty, Tom Clancy's Rainbow Six)
 Simulation of life (ex: SimCity, the SIMS)
 Simulation of vehicles (ex: Flight Simulators, Sailing Simulators)
 Racing games (ex: Nascar, Grand Prix)
 Sports games (ex: Football, Baseball)
 Other, please describe: _____
 I don't play videogames

What is your age? _____

What is your sex?

- Male Female

Date: _____

Participant number: _____

Post-study Questionnaire

Describe your feelings and reactions as you were controlling the object during your first attempt.

Describe your feelings and reactions as you were controlling the object during your second attempt.

Do you feel like you learned how to maneuver the object during the second attempt to the simulation?

Describe your feelings and reactions as you were controlling the object during your third attempt.

Do you feel like your performance was improved over the first performance? If so, rate on a scale of 1 to 5 your improvement, with 1 indicating a small improvement and 5 indicating a dramatic one.

Please indicate any comments, reactions or questions that you may have.

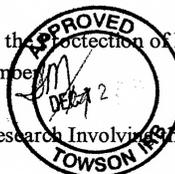
Date: _____

Control – Experimental

Participant number: _____

**APPROVAL NUMBER: 07-A045**

To: Giovanni Vincenti
From: Institutional Review Board for the Protection of Human
Subjects, Stephen Mogge, Member
Date: Tuesday, December 12, 2006
RE: Application for Approval of Research Involving the Use of
Human Participants



Office of University
Research Services

Towson University
8000 York Road
Towson, MD 21252-0001
t. 410 704-2236
f. 410 704-4494

Thank you for submitting an Application for Approval of Research Involving the Use of Human Participants to the Institutional Review Board for the Protection of Human Participants (IRB) at Towson University. The IRB hereby approves your proposal titled:

Fuzzy mediation as an improved ,method towards motor learning

If you should encounter any new risks, reactions, or injuries while conducting your research, please notify the IRB. Should your research extend beyond one year in duration, or should there be substantive changes in your research protocol, you will need to submit another application for approval at that time.

We wish you every success in your research project. If you have any questions, please call me at (410) 704-2236.

CC: ✓ G, Trajkovski
File



Date: Tuesday, December 12, 2006

NOTICE OF APPROVAL

TO: Giovanni Vincenti **DEPT:** COSC

PROJECT TITLE: *Fuzzy mediation as an improved method towards motor learning*

SPONSORING AGENCY: GSA

APPROVAL NUMBER: 07-A045

The Institutional Review Board for the Protection of Human Participants has approved the project described above. Approval was based on the descriptive material and procedures you submitted for review. Should any changes be made in your procedures, or if you should encounter any new risks, reactions, injuries, or deaths of persons as participants, you must notify the Board.

A consent form: is is not required of each participant

Assent: is is not required of each participant

This protocol was first approved on: 12-Dec-2006

This research will be reviewed every year from the date of first approval.



Stephen Mogge, Member
Towson University Institutional Review Board

Appendix B – Responses to Pre- and Post-Study Questionnaires

(Experiment 1)

Pre-study Questionnaires

Participant ID: 001/H1

Question 1: No
 Question 2: n/a
 Question 3: Yes
 Question 4: 1
 Question 5: Classroom learning, very basic
 Question 6: 8
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Mouse
 Question 10: Action, Precision shooting, Simulation of vehicles, Racing games, Sports games
 Age: 21
 Sex: Female

Participant ID: 002/H1

Question 1: No
 Question 2: n/a
 Question 3: Yes
 Question 4: 3
 Question 5: Remote control boats and airplanes, Flight simulator software
 Question 6: 8
 Question 7: 3-5 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Mouse
 Question 10: Action (most), Strategy
 Age: 35
 Sex: Male

Participant ID: 003/H1

Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 7
 Question 7: 3-5 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Mouse, Joystick
 Question 10: Strategy, Sports games
 Age: 22
 Sex: Male

Participant ID: 004/H1

Question 1: No
 Question 2: n/a
 Question 3: Yes
 Question 4: 7
 Question 5: Some experience driving boats
 Question 6: 9
 Question 7: 3-5 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Mouse
 Question 10: Action, Precision shooting, Racing, Sports
 Age: 21
 Sex: Male

Participant ID: 005/H1

Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: 1
 Question 5: n/a
 Question 6: 2
 Question 7: 1-2 times a week
 Question 8: Read the instructions manual
 Question 9: Mouse
 Question 10: I don't play videogames
 Age: 21
 Sex: Female

Participant ID: 006/H1

Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 7
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Mouse
 Question 10: Strategy
 Age: 35
 Sex: Male

Participant ID: 007/H1

Question 1: Yes
 Question 2: Small boats, Ace combat, Flight Sim
 Question 3: Yes
 Question 4: 5
 Question 5: Video games
 Question 6: 10
 Question 7: 3-5 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Mouse, Joystick
 Question 10: Action, Strategy, Simulation of vehicles
 Age: 27
 Sex: Male

Participant ID: 008/H1

Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 8
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Joystick
 Question 10: Action, Strategy, Simulation of life, Racing
 Age: 23
 Sex: Male

Participant ID: 009/H1
 Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 7
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Mouse, Joystick
 Question 10: Action, Precision shooting
 Age: 32
 Sex: Male

Participant ID: 010/H1
 Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 9
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Mouse
 Question 10: Strategy, Sports games
 Age: 31
 Sex: Male

Participant ID: 011/H1
 Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 4
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Joystick
 Question 10: Action, Sports
 Age: 18
 Sex: Male

Participant ID: 012/H1
 Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: n/a
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Joystick
 Question 10: Strategy
 Age: 27
 Sex: Male

Participant ID: 013/H1
 Question 1: Yes
 Question 2: Submarines (SSBN)
 Question 3: Yes
 Question 4: 10
 Question 5: I was a US Navy Submarine for five years
 Question 6: 9
 Question 7: 3-5 times a week
 Question 8: Read the instructions manual
 Question 9: Joystick
 Question 10: Simulation of life, Racing, Sports
 Age: 25
 Sex: Male

Participant ID: 014/H1
 Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 7
 Question 7: n/a
 Question 8: n/a
 Question 9: n/a
 Question 10: I don't play videogames
 Age: 18
 Sex: Female

Participant ID: 015/H1
 Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 6
 Question 7: 6-10 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Mouse
 Question 10: Action, Simulation of life, Second life
 Age: 25
 Sex: Male

Participant ID: 016/H1
 Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: 1
 Question 5: I have been on an airplane and boat before
 Question 6: 7
 Question 7: 1-2 times a week
 Question 8: Read the instructions manual
 Question 9: Keyboard, Mouse, Joystick
 Question 10: Action, Precision shooting, Sports
 Age: 19
 Sex: Male

Participant ID: 017/H1
 Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 9
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Joystick
 Question 10: Action, Sports
 Age: 18
 Sex: Male

Participant ID: 018/H1
 Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 8
 Question 7: n/a
 Question 8: Start playing and figure things out as you go
 Question 9: n/a
 Question 10: I don't play videogames
 Age: 19
 Sex: Female

Participant ID: 019/H1
 Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 6
 Question 7: 3-5 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Joystick
 Question 10: Sports
 Age: 18
 Sex: Male

Participant ID: 020/H1
 Question 1: No
 Question 2: n/a
 Question 3: Yes
 Question 4: 4
 Question 5: Experience commandeering boats in the ocean
 Question 6: 7
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Mouse, Joystick
 Question 10: Action, Precision shooting, Racing
 Age: 18
 Sex: Male

Participant ID: 021/H1
 Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 7
 Question 7: 6-10 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Mouse
 Question 10: Action, Strategy, Precision shooting
 Age: 18
 Sex: Male

Post-study Questionnaires

Participant ID: 001/H1
 Question 1: I thought it was fast and hard to control. It was over really quickly
 Question 2: I was more in control and was more successful.
 Question 3: I was a lot better the second time around.
 Question 4: I felt comfortable with it. I was happy to get all the dots.
 Question 5: 5. No dots on the first attempt. All dots on the third attempt.
 Question 6: The more I tried, the better I got. I felt more prepared each time.

Participant ID: 002/H1
 Question 1: Frustrated with the controls for the first few seconds
 Question 2: Had more confidence, but still frustrated that I missed the first dot
 Question 3: Yes
 Question 4: Confident but surprised that I still missed the first target
 Question 5: 3
 Question 6: Do I get paid for this? (Open bar should be employed in future trials)

Participant ID: 003/H1
 Question 1: Had to figure out quickly which was to move the joystick to go left and right but about halfway through it was pretty easy
 Question 2: Had to move against a force not shown on the screen
 Question 3: It was a little steadier in controlling it but I don't think I learned anything new from the first attempt
 Question 4: It seemed harder to control than the previous 2 attempts
 Question 5: 1
 Question 6: n/a

Participant ID: 004/H1
 Question 1: Frustrated that I could not get control of the controls as quickly as I expected
 Question 2: Though I was doing better than I actually was
 Question 3: Somewhat but I still did not have a full grasp of how sensitive the device was
 Question 4: Began to feel somewhat comfortable with the interface
 Question 5: Yes, I feel my performance did improve as by the last time I felt like I better understand the controls and the sensitivity of the device. I feel my improvement was a 3 on the 1-5 scale
 Question 6: n/a

Participant ID: 005/H1
 Question 1: Anticipation, anxious feelings, difficulty controlling object
 Question 2: Better control, able to anticipate motion
 Question 3: Yes
 Question 4: I had poor control because I have little video game experience
 Question 5: 2
 Question 6: n/a

Participant ID: 006/H1
 Question 1: n/a
 Question 2: n/a
 Question 3: n/a
 Question 4: n/a
 Question 5: n/a
 Question 6: n/a

Participant ID: 007/H1

Question 1: Surprised by speed
 Question 2: Easy to control
 Question 3: No at the end of 1
 Question 4: Focused but easy
 Question 5: 3
 Question 6: n/a

Participant ID: 008/H1

Question 1: A bit confused and rushed as I did not know what to do the first 2 seconds
 Question 2: Relaxed, calm and confident on what to do and expect
 Question 3: Yes I do
 Question 4: Definitely relaxed and confident as I knew what to expect and what was expected from me
 Question 5: 4, my performance definitely improved
 Question 6: n/a

Participant ID: 009/H1

Question 1: Was trying to figure out the way the object reacted to the joystick movement-sensitivity
 Question 2: More relaxed and in control
 Question 3: Yes
 Question 4: More confident
 Question 5: Yes, 4
 Question 6: n/a

Participant ID: 010/H1

Question 1: Caught a little off guard to start
 Question 2: Easy – trusted the “co-pilot” and myself. Just focused on navigating. Pretended copilot was not there
 Question 3: No
 Question 4: Easier
 Question 5: 2
 Question 6: Might be too simple to really observe or experience improvement

Participant ID: 011/H1

Question 1: I was surprised at first when I realized the controls are reverse
 Question 2: I was frustrated the object was moving so slow
 Question 3: I felt like I controlled it about the same during every attempt
 Question 4: I was thinking about getting the object centered on the landing strip
 Question 5: Yes, 2
 Question 6: n/a

Participant ID: 012/H1

Question 1: It was fast I had to get adjusted to the speed of the game
 Question 2: I was a lot slower. Easier to control.
 Question 3: A little, but I did not feel as I had it down packed
 Question 4: It was still fast hard to control but I was ready and more adjusted this time
 Question 5: Yes, 2
 Question 6: n/a

Participant ID: 013/H1

Question 1: It's just started immediately, I was caught a little off guard by the sensitivity.
 Question 2: It was much slower and I really felt more in control. I was also more aware of the sensitivity of the joystick.
 Question 3: Yes
 Question 4: I knew that it would be fast but I also understood the sensitivity so it made it easier than the first time. I wasn't caught off guard...
 Question 5: 4
 Question 6: n/a

Participant ID: 014/H1

Question 1: I was a little confused as to what I was doing.
 Question 2: It was a lot less confusing because it was slower.
 Question 3: I felt that I learned a little during the second attempt to the simulation.
 Question 4: I felt like I understood it more
 Question 5: I would give myself a 1
 Question 6: n/a

Participant ID: 015/H1

Question 1: It was hard to get adjusted the first time
 Question 2: It seemed much easier to hit the circles and the flight path was much more accurate
 Question 3: Yes it was helpful
 Question 4: I felt that it was easier to hit the objects + I felt that I understood how to control the aircraft better
 Question 5: Yes, 1 for improvement
 Question 6: Great program. The flight path relative to the screen direction was interesting.

Participant ID: 016/H1

Question 1: The first attempt was the dardest. I was moving the joystick around too much and the plane was flying really fast.
 Question 2: At the slower speed I felt a lot more comfortable and I got the feel of the joystick, so I hit every point before landing.
 Question 3: Yes, I definitely learned how to maneuver the plane by just slightly moving the joystick during the second attempt.
 Question 4: It was smooth and easier to control.
 Question 5: 5, because on my last run I nailed all the points before landing and on my first run I missed one.
 Question 6: I liked using the joystick like I was flying a plane.

Participant ID: 017/H1

Question 1: I was quite sure what to expect but after controlling at the end it OK.
 Question 2: I felt more in control because I saw the course and knew what to expect and was more similar with the controls.
 Question 3: Yes
 Question 4: I felt that I was in complete control when going at normal speed + going through it twice I was able to perform well
 Question 5: Yes I feel my performance improved from hitting only 1 or 2 checkpoints in the beginning to hitting all on the third try.
 Question 6: Controlling was a little weird but after the first try it became simple.

Participant ID: 018/H1

Question 1: I was confused why when it started, I was moving the joystick right but the object I was trying to move was going left.
 Question 2: I realized that when the object was heading down the screen, I needed to move the joystick left to make the object move right.
 Question 3: Yes
 Question 4: It was still harder than the second time [because] of the speed, but I understood what to do better.
 Question 5: 2 – 3
 Question 6: n/a

Participant ID: 019/H1

Question 1: I was confused at first on how the plane was going to move and then found that if I turned the joystick to the right the plane moved right like normal.
 Question 2: I felt confident and in control since I already had experience with it.
 Question 3: Yes
 Question 4: I felt in control and confident that I would do much better than my first attempt.
 Question 5: Yes. Improvement equalling: 4.
 Question 6: n/a

Participant ID: 020/H1

Question 1: The object was moving very fast and I had to move it around a few times to get a feel for it.

Question 2: The object was moving much slower and was much easier to control

Question 3: Yes, I got a feel for how it reacted although it was very different moving slower.

Question 4: Similar to my first attempt although I got more of a feel for it.

Question 5: 2 – 3

Question 6: n/a

Participant ID: 021/H1

Question 1: Panic, didn't know left and right.

Question 2: Easier [because] it was slower, but a little intimidating doing a co-pilot.

Question 3: Yes

Question 4: A little more control but still difficult [because] it was fast

Question 5: 3

Question 6: n/a

Appendix C – Responses to Pre- and Post-Study Questionnaires

(Experiment 2)

Pre-Study Questionnaires

Participant ID: 001/H2

Question 1: Yes
 Question 2: Submarines (SSBN)
 Question 3: Yes
 Question 4: 10
 Question 5: I was a US Navy Submarine for five years
 Question 6: 9
 Question 7: 3-5 times a week
 Question 8: Read the instructions manual
 Question 9: Joystick
 Question 10: Simulation of life, Racing, Sports
 Age: 25
 Sex: Male

Participant ID: 002/H2

Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 7
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Mouse
 Question 10: Strategy
 Age: 35
 Sex: Male

Participant ID: 003/H2

Question 1: No
 Question 2: n/a
 Question 3: Yes
 Question 4: 5
 Question 5: Flight sim
 Question 6: 8
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Joystick
 Question 10: Simulation of vehicles
 Age: 28
 Sex: Male

Participant ID: 004/H2

Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 6
 Question 7: 6-10 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Mouse
 Question 10: Action, Simulation of life, Second life
 Age: 25
 Sex: Male

Participant ID: 005/H2

Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 9
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Joystick
 Question 10: Action, Sports
 Age: 18
 Sex: Male

Participant ID: 006/H2

Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: 1
 Question 5: I have been on an airplane and boat before
 Question 6: 7
 Question 7: 1-2 times a week
 Question 8: Read the instructions manual
 Question 9: Keyboard, Mouse, Joystick
 Question 10: Action, Precision shooting, Sports
 Age: 19
 Sex: Male

Participant ID: 007/H2

Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 8
 Question 7: n/a
 Question 8: Start playing and figure things out as you go
 Question 9: n/a
 Question 10: I don't play videogames
 Age: 19
 Sex: Female

Participant ID: 008/H2

Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 6
 Question 7: 3-5 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Joystick
 Question 10: Sports
 Age: 18
 Sex: Male

Participant ID: 009/H2
 Question 1: No
 Question 2: n/a
 Question 3: Yes
 Question 4: 4
 Question 5: Experience commandeering boats in the ocean
 Question 6: 7
 Question 7: 1-2 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Mouse, Joystick
 Question 10: Action, Precision shooting, Racing
 Age: 18
 Sex: Male

Participant ID: 010/H2
 Question 1: No
 Question 2: n/a
 Question 3: No
 Question 4: n/a
 Question 5: n/a
 Question 6: 7
 Question 7: 6-10 times a week
 Question 8: Start playing and figure things out as you go
 Question 9: Keyboard, Mouse
 Question 10: Action, Strategy, Precision shooting
 Age: 18
 Sex: Male

Post-Study Questionnaires

Participant ID: 001/H2
 Question 1: I felt really calm and confident. I understood the sensitivity and was in control.
 Question 2: I knew how sensitive it would be, however I was unsure about how my left hand would perform. But I knew once I got passed the first circle I would have control.
 Question 3: Yes
 Question 4: I felt like my right hand was determining which way I would maneuver more than the left hand.
 Question 5: No
 Question 6: n/a

Participant ID: 002/H2
 Question 1: It was easier as I was trying to mimic the right hand commands from previous trials
 Question 2: Frustrating. Focus was right hand.
 Question 3: Yes
 Question 4: Much easier
 Question 5: Yes, 4
 Question 6: n/a

Participant ID: 003/H2
 Question 1: Felt good about control, still unsure though
 Question 2: Good about a somewhat steady green indicator
 Question 3: If anything, I was more confident
 Question 4: Much more confident
 Question 5: Yes, 3
 Question 6: n/a

Participant ID: 004/H2
 Question 1: Using my left hand felt more awkward then using my right hand
 Question 2: It seemed difficult at first to coordinate both hands but it was easier towards the end
 Question 3: Yes, after flying a few times I felt more confident
 Question 4: It was easier the third time
 Question 5: Yes I think there was an improvement, 3
 Question 6: n/a

Participant ID: 005/H2
 Question 1: I felt under control because I had already been through the course rightly
 Question 2: Once again I felt completely in control
 Question 3: I felt that I was in complete control of the object
 Question 4: Again in complete control having gone through the course five times already
 Question 5: Yes and I feel I improved to a 5
 Question 6: After a few rounds with the game it was quite simple

Participant ID: 006/H2
 Question 1: The left hand felt just as comfortable as my right
 Question 2: The speed wasn't a factor I could fly the plane with my eyes [closed].
 Question 3: No, I learned it with my right hand left came natural.
 Question 4: It was smooth sailing
 Question 5: 2, because I nailed all the points everytime on each run before landing.
 Question 6: I was equally accurate with both hands.

Participant ID: 007/H2
 Question 1: I was scared to use my left hand, but it seemed easier than when I used my right hand
 Question 2: It was still easier than when I used my right hand. I was confused by this.
 Question 3: Yes, but also during the first attempt.
 Question 4: Still easier than [with] my right hand. I wondered if the speed increased, which might have distracted me.
 Question 5: 3 or 4
 Question 6: n/a

Participant ID: 008/H2
 Question 1: It was different [due] to the fact that ... it was my weaker hand and I felt very out of control in attempting to control the plane.
 Question 2: I felt more confident and in more control.
 Question 3: Yes
 Question 4: I felt really confident that I was in more control with the plane
 Question 5: Yes. 3.
 Question 6: n/a

Participant ID: 009/H2
 Question 1: Didn't move too fast and was easy to control
 Question 2: Easier to move, but not by too much, not too much improvement
 Question 3: Learned how to move the object a little more but not my too much
 Question 4: Very easy to move around, took little skill or concentration
 Question 5: 1
 Question 6: n/a

Participant ID: 010/H2

Question 1: I thought I was going to do bad but because it was slow it was easy.

Question 2: A lot more control.

Question 3: Yes

Question 4: Knew how to do it.

Question 5: 2

Question 6: n/a

List of References

- Abad, P., Suárez, A., Gasca, R., & Ortega, J. (2002). Using Supervised Learning Techniques for Diagnosis of Dynamic Systems. *Proceedings of the International Workshop on Principles of Diagnosis*.
- Agrawal, R., Imielinski, T. & Swami, A. (1993). Mining Association Rules Between Sets of Items in Large Databases. *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data*, 207-216.
- Arai, S. & Ishida T. (2004). Learning for Human-Agent Collaboration on the Semantic Web. *Proceedings of the 12th International Conference on Informatics Research for Development of Knowledge Society Infrastructure*, 132-139.
- Barnes C., Baranyi A., Bindman L., Dudal Y., Fregnac Y., Ito M., Knopfel T., Lisberger S., Morris R., Moulins M., Movshon J., Singer W. & Squirre L. (1994). Group report: Relating activity-dependent modifications of neuronal function to changes in neural systems and behavior. In A. I. Selverston and P. Ascher (Eds), *Cellular and Molecular Mechanisms Underlying Higher Neural Functions*. Wiley, New York.
- Bennett, C. & Pilkington, R. (2001). Using a Virtual Learning Environment in Higher Education to Support Independent and Collaborative Learning. *Proceedings of the IEEE International Conference on Advanced Learning Technologies*, 285-288.
- Borenstein, J. & Ulrich, I. (1997). The GuideCane – A Computerized Travel Aid for the Active Guidance of Blind Pedestrians. *Proceedings of the IEEE International Conference on Robotics and Automation*, 1283-1288.
- Carroll, J., & Mazur, S. (1986). Lisa Learning. *IEEE Computer*, 19(11), 35-49.
- Caterpillar. (n.d.). Virtual Training Solutions. Retrieved August 2, 2007, from <http://www.cat.com/cda/layout?m=96642&x=7>.
- Chialvo, D. & Bak P. (1999). Learning from Mistakes. *Neuroscience*, 90(4), 1137-1148.
- Code of Federal Regulations (July 31, 2007). Title 14: Aeronautics and Space.
- Colgate, J., Wannasuphprasit, W. & Peshkin, M. (1996). Cobots: Robots for Collaboration with Human Operators. *International Mechanical Engineering Congress and Exposition*.
- Connell, J. & Viola, P. (1990). Cooperative Control of a Semi-Autonomous Mobile Robot. *Proceedings of the 1990 IEEE International Conference on Robotics and Automation*, 1118-1121.
- Dugan, B. & Glinert, E. (2002). Task Division in Collaborative Simulations. *Proceedings of the 35th Annual Hawaii International Conference on System Sciences (HICSS-35'02)*, 3.1.
- Fiaschetti, A. (Personal communication, April 30, 2007).
- Gokhale, A. (1995). Collaborative Learning Enhances Critical Thinking. *Journal of Technology Education*, 7(1), 22-30.

- Gomi, T. & Griffith, A. (1998). Developing Intelligent Wheelchairs for the Handicapped. *Lecture Notes In Computer Science*, 1458, 150-178.
- Goncharenko, I., Svinin, M., Matsumoto, S., Masui, Y., Kanou, Y. & Hosoe, S. (2004). Cooperative Control with Haptic Visualization in Shared Virtual Environments. *Proceedings of the 8th International Conference on Information Visualisation*, 533-538.
- Grenoble O'Malley, P. (2005, March/April). Troubleshooting Machine Control: Contractors Talk About Making the Most of 3D GPS. *Grading & Excavator Contractor*, 7(2).
- Griffin, W. (2003). *Shared Control For Dexterous Telematipulation With Haptic Feedback*. Ph.D. Thesis.
- Griffin, W., Provancher, W. & Cutkosky, M. (2005). Feedback Strategies for Telematipulation with Shared Control of Object Handling Forces. *Presence: Teleoperators and Virtual Environments*, 14(6), 720-731.
- Griffiths, P. & Gillespie, R. (2004). Shared Control Between Human and Machine: Haptic Display of Automation During Manual Control of Vehicle Heading. *Proceedings of the 12th International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 358-366.
- Griffiths, P. & Gillespie, R. (2005). Sharing Control Between Human and Automation Using Haptic Interface: Primary and Secondary Task Performance Benefits. *Human Factors*.
- Hammonds, K., Jackson S., DeGeorge, G. & Morris, K. (1997). The New University – A Tough Market is Reshaping Colleges. *Business Week*, 22, 96-102.
- Hebb, D. (1949). *The Organization of Behavior*. New York: John Wiley.
- Heyes, C. (2001). Causes and consequences of imitation. *TRENDS in Cognitive Sciences*, 5(6), 253-261.
- Hopfield, J. (1982). Neural Networks and Physical Systems with Emergent Collective Computational Abilities. *Proceedings of the National Academy of Sciences*, 79, 2554-2558.
- ISA. (May 31, 2006). Airbus expands cockpit automation. Retrieved August 2 2007 from <http://www.isa.org/InTechTemplate.cfm?Section=Business1&template=/ContentManagement/ContentDisplay.cfm&ContentID=54287>
- Kandel, E., Schwartz, J. & Jessell, T. (2000). *Principles of Neural Sciences*. New York: The McGraw-Hill Companies.
- Karamanoukian, H., Pande, R., Patel, Y., Freeman, A., Aoukar, P. & D'Ancona, G. (2003). *Pediatric Endosurgery & Innovative Techniques*, 7(4), 421-425.
- Kokar, M., Tomasik, J. & Weyman, J. (2004). Formalizing Classes of Information Fusion Systems. *Information Fusion*, 5(3), 189-202.
- Kwek, S. (1999). An Apprentice Learning Model. *Proceedings of the twelfth annual conference on Computational learning theory*, 63-74.
- Levine, S., Bell, D., Jaros, L., Simpson, R. & Koren, Y. (1999). The NavChair Assistive

- Wheelchair Navigation System. *Proceedings of the 1995 International Joint Conference on Artificial Intelligence*.
- Lexus. (n.d.). Advanced Parking Guidance System. Retrieved on August 2, 2007 from http://www.lexus.com/models/LS/features/exterior/advanced_parking_guidance_system.html.
- Li, Y., Patoglu, V. & O'Malley, M. (2006). Shared Control for Training in Virtual Environments: Learning Through Demonstration? *Proceedings of Eurohaptics*.
- Liu, S. & Silverman, M. (2001). A practical guide to biometric security technology. *IT Professional*, 3(1), 27-32.
- Mamdani, E. & Assilian S. (1975) An experiment in linguistic synthesis with a fuzzy logic controller. *International Journal of Man-Machine Studies*, 7(1), 1-13.
- Martini, F. (2005). *Fundamentals of Anatomy & Physiology*. San Francisco, CA: Benjamin Cummings.
- Mendel, J. (2003). Type-2 Fuzzy Sets: Some Questions Answered. *Newsletter of the IEEE Neural Networks Society*, 1, 10-13.
- Mendel, J. & John, R. (2002). Type-2 Fuzzy Sets Made Simple. *IEEE Transactions on Fuzzy Systems*, 10(2), 117-127.
- Niccolai, J. (March 12, 2004). Robot Cars Race in the Desert. Retrieved on August 2, 2007 from <http://www.pcworld.com/article/id,115198-page,1/article.html>.
- Noguchi, N., Reid, J., Benson, E., Will, J. & Stombaugh, T. (1998). Vehicle automation system based on multisensor integration. *ASAE Paper 983111*.
- Norman, D. (2005, November-December). There's an Automobile in HCI's Future. *Interactions*, 53-54.
- O'Malley, M., Gupta, A., Gen, M. & Li, Y. (2006). Shared Control in Haptic Systems for Performance Enhancement and Training. *Journal of Dynamic Systems, Measurement, and Control*, 128(1), 75-85.
- Piaget, J. (1945). *Play, Dreams, and Imitation in Childhood*. New York: Norton.
- Pedrycz, W. & Gomide, F. (1998). *An Introduction to Fuzzy Sets: Analysis and Design*. Cambridge, MA: A Bradford Book.
- Reid, J. (1998). *Precision Guidance of Agricultural Vehicles*. Retrieved January 2007 from <http://www.age.uiuc.edu/oree/pdf/uilu-eng-98-7031.pdf>.
- Ribeiro Filho, J., Treleaven, P. & Alippi, C. (1994). Genetic-algorithm programming environments. *Computer*, 27(6), 28-43.
- Rizzolatti, G., Fadiga, L., Gallese, V. & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3(2) 131-141.
- Ross, A. & Jain, A. (2003) Information fusion in biometrics. *Lecture Notes in Computer Science*, 24(13), 2115-2125.
- Rosser, J., Lynch, P., Cuddihy, L., Gentile, D., Klonsky, J. & Merrell, R. (2007). The Impact of Video Games on Training Surgeons in the 21st Century. *Archives of Surgery*,

142, 181-186.

Rovira Más, F., Zhang, Q., Reid, J. & Will, J. (2005). Visual crop row detection algorithm for automated guidance. *Journal of Automobile Engineering*, 219(8), 999-1010.

Roy, N., Baltus, G., Fox, D., Gemperle, F., Goetz, J., Hirsch, T., Margaritis, D., Montemerlo, M., Pineau, J., Schulte, J. & Thrun, S. (2000). Towards personal service robots for the elderly. *Workshop on Interactive Robots and Entertainment*.

Shelley, G., Cashman, T. & Vermaat, M. (2007). *Discovering Computers*. Boston, MA: Thomson Course Technology.

Shinohara, T., Miyawaki, K. & Sunayama, K. (2001). Heavy material handling by cooperative robot control. *Advanced Robotics*, 15(3), 317-321.

Steele, M. & Gillespie, R. (2004). Shared Control Between Human and Machine: Haptic Display of Automation During Manual Control of Vehicle Heading. *Proceedings of the 12th Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems*, 358-366.

Sudkamp, T. & Hammell, R. (1994). Interpolation, Completion, and Learning Fuzzy Rules. *IEEE Transactions on Systems, Man, and Cybernetics*, 24(2), 332-342.

Thorndike, E. (1911) *Animal Intelligence*. New York: Macmillan.

Trajkovski, G. (2005). E-POPSICLE: An Online Environment for Studying Context Learning in Human and Artificial Agents. *Proceedings of the 16th Midwest AI and Cognitive Science Conference*, 61-66.

Trajkovski, G. (2007). *An Imitation-Based Approach to Modeling Homogenous Agents Societies*. Hershey, PA: IDEA Publishing.

Trajkovski, G., Schlosburg, J., Whitman, B. & Vincenti, G. (2005). Building infrastructure for an honors research robotics lab. *Proceedings of the IEEE 6th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing*, 352-357.

Trajkovski, G., Stojanov, G. & Vincenti, G. (2005). Extending MASIVE: The impact of stress on imitation-based learning. *Proceedings of the IEEE 6th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing*, 360-365.

Tran, C., Jain, L. & Abraham, A. (2002) TACDSS: Adaptation of a Takagi–Sugeno Hybrid Neuro-Fuzzy System. *Proceedings of the 7th Online World Conference on Soft Computing in Industrial Applications*.

Tsugawa, S., Kato, S., Tokuda, K., Matsui, T. & Fujii, H. (2001). A cooperative driving system with automated vehicles and inter-vehicle communications in Demo 2000. *Intelligent Transportation Systems*, 918 - 923.

Vincenti, G., Hammell, R. & Trajkovski, G. (2005). Data Mining for Imprecise Temporal Associations. *Proceedings of the IEEE 6th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing*, 76-

81.

- Vincenti G. & Trajkovski, G. (2006). Fuzzy Mediation for Online Learning in Autonomous Agents. *2006 Fall AAI Symposium*, 127-133.
- Vincenti G. & Trajkovski, G. (2007). Fuzzy Mediation as a Dynamic Extension to Information Fusion. *Fusion 2007*.
- Wallace, J. (2000, March 20). Unlike Airbus, Boeing lets aviator override fly-by-wire technology. *Seattle Post-Intelligencer*.
- Wang, L. & Mendel, J. (1992). Generating Fuzzy Rules by Learning from Examples. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(6): 1414-1427.
- Wasson, G. & Gunderson, J. (2001). Variable autonomy in a Shared Control Pedestrian Mobility Aid for the Elderly. *Proceedings of the IJCAI'01 Workshop on Autonomy, Delegation, and Control*.
- Wasson, G., Gunderson, J., Graves, S. & Felder, R. (2001). Effective Shared Control in Cooperative Mobility Aids. *Proceedings of the 14th International Florida Artificial Intelligence Research Society Conference*, 509-513.
- Wasson, G., Sheth, P., Alwan, M., Granata, K., Ledoux, A. & Huang, C. (2003). User Intent in a Shared Control Framework for Pedestrian Mobility Aids. *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- WordNET Search (n.a.). *Control*. Retrieved February 16th, 2006 from <http://wordnet.princeton.edu/perl/webwn?s=control>.
- Yang, S., Park, S., Jeong, J. & Park, B. (1999). Development of intelligent automated driving control system(lateral control). *Proceedings of the 3rd Russian-Korean International Symposium on Science and Technology*, 1, 334-337.
- Yanko, H. (1998). Wheelchair: A Robotic Wheelchair System: Indoor Navigation and User Interface. *Lecture Notes In Computer Science*, 1458, 256-268.
- Yen, J. & Langari, R. (1999). *Fuzzy Logic: Intelligence, Control and Information*. Upper Saddle River, NJ: Prentice-Hall.
- Ysander, L. (1970). Sick and handicapped drivers. A study on the risks of sudden illness at the wheel and on the frequency of road accidents and traffic offenses in chronically sick, disabled, and elderly drivers. *ACTA Chirurgica Scandinavica – Supplementum*, 409, 1-27.
- You, S., Jang, S., Kim, Y., Hallett, M., Ahn, S., Kwon, Y., Kim, J. & Lee, M. (2005). Virtual Reality-Induced Cortical Reorganization and Associated Locomotor Recovery in Chronic Stroke: An Experimenter-Blind Randomized Study. *Stroke*, 36, 1166-1171.
- Zadeh, L. (1965). Fuzzy Sets. *Information and Control*, 8, 338-353.
- Zhang, B., Chen, Y., Fan, W., Fox, E., Gonçalves, M., Cristo, M. & Calado, P. (2005). Intelligent GP Fusion from Multiple Sources for Text Classification. *Proceedings of the 14th ACM international conference on Information and knowledge management*.

Curriculum Vitae

Giovanni Vincenti

giovanni.vincenti@gmail.com

Via dei Gracchi, 123
Rome, Italy
00192
+39 (06) 32.43.268

Education

Towson University **Bachelor of Arts in Biology – Received in May 2001**
Towson, MD 21252, Concentration in Human Anatomy and Physiology (Pre-Medicine)
USA

Towson University **Master of Science in Computer Sciences – Received in August 2004**
Towson, MD 21252, Concentration in Software Engineering
USA Masters Thesis: Data Mining for Imprecise Temporal Associations

Towson University **Doctorate of Science in Applied Information Technology – Expected August 2007**
Towson, MD 21252,
USA Area of research: Fuzzy mediation as an improved method towards machine learning and information fusion.

Appointments

05/2006 to Present **Adjunct Professor – Computer and Information Sciences Department**
Towson University
Towson, MD 21252, In charge of teaching two undergraduate-level courses:
USA

- Information and Technology for Business;
- Project Management.

Participate in research activities spanning over various domains:

- Concepts of A.I. and human cognition applied to the creation of new learning models for humans;
- Data mining techniques applied to medical informatics;
- Creation of a robotics infrastructure that will accommodate learning as well as research.

08/2004 to 05/2006 **Research Assistant – Computer and Information Sciences Department**
Towson University
Towson, MD 21252, Participate in research activities spanning over various domains:
USA

- Concepts of A.I. and human cognition applied to the creation of new learning models for humans;
- Data mining techniques applied to medical informatics;
- Image processing and multiagency theory applied to telerobotics;
- Creation of a robotics infrastructure that will accommodate learning as well as research.

In charge of teaching two undergraduate-level courses:

- Information and Technology for Business;
- Data Organization.

8/2003 to 8/2004 **Graduate Assistant, Webmaster – Faculty Development**

Towson University

Towson, MD 21252,

USA

Managed various websites that were linked to the department:

- Responsible for the design and creation two new websites;
- Support the original departmental website.

Organized several aspects of two Lilly Conferences on College and University Teaching (East):

- Responsible for the creation of the printed material;
- Curated most of the graphic design aspects;
- Management of the proposals.

8/2002 to 8/2003 **Graduate Assistant, Programmer and Webmaster – Child Daycare Center**

Towson University

Towson, MD 21252,

USA

Responsible for the following tasks:

- Maintenance and redesign of the center's website;
- Development of a program that managed student-observers.

8/2001 to 8/2002 **Lab Assistant – Office of Technical Support**

Towson University

Towson, MD 21252,

USA

Performed various technical support duties including:

- Supervision of 120 computers dedicated to the students;
- Technical support for questions about the software installed on the computers.

8/1999 to 12/2001 **Student Mentor – Biology Department**

Towson University

Towson, MD 21252,

USA

Responsible for the following activities:

- Taught concepts of biology, chemistry and physics to students in small groups;
- Tutored any students who needed help on understanding peer-reviewed journal articles;
- Helped in the design of improvements for the course being taught.

Publications

2007

Vincenti, G. & Trajkovski, G. (2007). Analysis of Different Mediation Equations and Tightness of Control to Finely Regulate the Exchange of Control Between Expert and Novice Controllers in a Fuzzy Mediation Environment. *NAFIPS 2007*. June 24-27, 2007. San Diego, CA. USA.

Vincenti, G. & Trajkovski, G. (2007). Fuzzy Mediation as a Dynamic Extension to Information Fusion Applied to Learning Environments. *11th International Conference on Cognitive and Neural Systems*. May 16-19, 2007. Boston, MA. USA.

2006

Vincenti G. and Trajkovski, G. (2006). Fuzzy Mediation for Online Learning in Autonomous Agents. *2006 Fall AAAI Symposium*. October 12-15, 2006. Arlington, VA. USA.

Vincenti G. and Trajkovski, G. (2006). Creating an Elastic Robotics Framework for Education and Research. *2006 Fall AAAI Symposium*. October 12-15, 2006. Arlington, VA. USA.

2005

Vincenti, G., Hammell, R. & Trajkovski, G. (2005). Scouting for Imprecise Temporal Associations to Support Effectiveness of Drugs During Clinical Trials. *Soft Computing for Real World Applications - NAFIPS 05*. June 22-25, 2005. Ann Arbor, MI. USA.

Trajkovski, G. & Vincenti, G. (2005). A Fuzzy Framework for Modelling Multiagent Societies. *Soft Computing for Real World Applications - NAFIPS 05*. June 22-25, 2005. Ann Arbor, MI. USA.

Vincenti, G., Hammell, R. & Trajkovski, G. (2005). Data Mining for Imprecise Temporal Associations. *6th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing – SNPD 2005*. May 23-25, 2005. Towson, MD. USA.

Trajkovski, G., Schlosburg, J., Whitman, B. & Vincenti, G. (2005). Building Infrastructure for an Honors Research Robotics Lab. *6th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing – SNPD 2005*. May 23-25, 2005. Towson University. Towson, MD. USA.

Trajkovski, G., Stojanov, G. & Vincenti, G. (2005). Extending MASIVE: The Impact of Stress on Imitation-Based Learning. *6th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing – SNPD 2005*. May 23-25, 2005. Towson, MD. USA.

2004

Vincenti G. & Hammell, R. (2004). Using Fuzzy Set Principles to Account for Imprecision in the Discovery of New Knowledge in Drug Analysis. *International Conference on Bioinformatics and its Applications – ICBA 04*. December 16-19, 2004. Fort Lauderdale, FL. USA.

Vincenti, G., Wilson, P., Crest, S. & Mangurian, L. (2004). Behind the Scenes in the 21st Century. An Interdisciplinary Approach Applied to Providing Information about Teaching and Learning for Faculty Members at a Comprehensive University. *7th Annual Lilly Conference on College and University Teaching, East*. April 2-3, 2004. Towson University. Towson, MD 21252.

Synergistic Activities

- | | |
|--------------------------------|---|
| <i>Conference Organization</i> | <ul style="list-style-type: none"> • Thirteenth Annual Multicultural Conference, Towson University, Towson, MD (March 2007) • 2006 AAAI Fall Symposium on the Interaction and emergent phenomena in societies of agents, Hyatt Crystal City, Arlington, VA (October 2006) • 6th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing, Towson University, Towson, MD (May 2005) • Lilly Conference on College and University Teaching – East 2005, Towson University, Towson, MD (April 2005) • Lilly Conference on College and University Teaching – East 2004, Towson University, Towson, MD (April 2004) |
| <i>Reviewer Activities</i> | <ul style="list-style-type: none"> • International Workshop on Fuzzy Logic and Applications. Genova (Italy). July 7-10, 2007. • Diversity in IT Education: Issues and Controversies. Book published in January 2006. |

- Professional Organization Memberships*
- President of the local chapter of Upsilon Pi Epsilon, International Honor Society for the Computing and Information Disciplines (December 2004 to May 2007)
 - Active member of Upsilon Pi Epsilon, International Honor Society for the Computing and Information Disciplines (since January 2004)

- Grants and Awards*
- Recipient of the Graduate Student Association award, Towson University (October 2006)
 - Recipient of the Graduate Student Fellowship, College of Graduate Studies and Research, Towson University (Academic year 2006-2007)
 - Recipient of the Graduate Student Association award, Towson University (April 2006)
 - Recipient of the UPE Award for Excellence and Service in Computer Science, Towson University (April 2006)
 - Recipient of the Graduate Student Association award, Towson University (December 2004)
 - Recipient of the NSF-funded University System of Maryland “Vertically Integrated Partnership K-16” Grant, Towson University (May 2003 to May 2004)