

9-2011

Variations on Stigmergic Communication to Improve Artificial Intelligence and Biological Modeling

Megan Marie Olsen
University of Massachusetts Amherst

Follow this and additional works at: https://scholarworks.umass.edu/open_access_dissertations



Part of the [Computer Sciences Commons](#)

Recommended Citation

Olsen, Megan Marie, "Variations on Stigmergic Communication to Improve Artificial Intelligence and Biological Modeling" (2011). *Open Access Dissertations*. 477.
<https://doi.org/10.7275/2395707> https://scholarworks.umass.edu/open_access_dissertations/477

This Open Access Dissertation is brought to you for free and open access by ScholarWorks@UMass Amherst. It has been accepted for inclusion in Open Access Dissertations by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

**VARIATIONS ON STIGMERGIC COMMUNICATION TO IMPROVE
ARTIFICIAL INTELLIGENCE AND BIOLOGICAL MODELING**

A Dissertation Presented

by

MEGAN M. OLSEN

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

September 2011

Computer Science

© Copyright by Megan M. Olsen 2011

All Rights Reserved

**VARIATIONS ON STIGMERGIC COMMUNICATION TO IMPROVE
ARTIFICIAL INTELLIGENCE AND BIOLOGICAL MODELING**

A Dissertation Presented

by

MEGAN M. OLSEN

Approved as to style and content by:

Hava T. Siegelmann, Chair

Ramesh Sitaraman, Member

Victor Lesser, Member

John Nambu, Member

Andrew Barto, Department Chair
Computer Science

ACKNOWLEDGEMENTS

One of the keys to success in modeling any non-computer system is to ensure that an expert is available to guide and analyze the system. I am therefore indebted to Nava Siegelmann-Danieli, Joe Jerry, and John Nambu for their guidance on the biological aspects of my work. I also appreciate the support of my advisor Hava Siegelmann, whose varied knowledge both within and outside of computer science, as well as her broad research interests, encouraged me to pursue many directions that are not traditional for a computer scientist. Additionally, I am very grateful to Ramesh Sitaraman and Victor Lesser, who have provided guidance and advice on different aspects of my research throughout my time at UMass. I appreciate the willingness of all of my committee members to be involved in my thesis, and to give feedback: Hava Siegelmann, Victor Lesser, Ramesh Sitaraman, and John Nambu.

I was very fortunate while at UMass to have met a variety of wonderful people, and make many friendships that helped ease the stress of graduate school. I am especially grateful to Yariv Levy for our Friday Lunches that helped keep us both sane and productive through many long weeks, and for his excellent feedback on every presentation. I am also thankful for the camaraderie of my other lab mates David Cooper, Kun Tu, Dimitri Nowicki, and Kyle Harrington, and their collaboration and feedback over the years. I will greatly miss their company, as well as that of Gal Niv, Filip Jagodzinski, Bobby Simidchieva, Stefan Christov, Jackie Feild, Henry Feild, Ilene Magpiong, Marc Cartright, Dirk Ruiken, Erin Cooper, Aruna Balasubramanian, and Niranjan Balasubramanian. I appreciate all of the laughter and good advice TJ Brunette gave me, as well as Sarah Osentoski's and Audrey St. John's constant willingness to give advice when asked. As there are not enough pages to list everyone who made a positive impact on my time in Amherst, I trust that they know who they are and that they are no less appreciated.

And of course, where would any of us be at this point without the helpfulness of all of our support staff? I am sincerely grateful to Leeanne Leclerc for always either knowing the answer or being able and willing to find the answer to my questions about requirements, rules, and regulations.

I am also grateful to both Priscilla Scott and Gwyn Mitchell for their administrative assistance to my lab, and Barbara Sutherland for her gracious assistance through the years. I also sincerely thank everyone in CSCF for all of the computer support that makes finishing a computer science PhD possible, and Andrew McCallum for providing me access to his computing cluster. Thank you also to the funding agencies that supported my research: Department of Homeland Security, National Science Foundation, and Office of Naval Research. Thank you to the many conferences that provided student travel funding, and to UMass Amherst and the Computer Science department for travel funding and TA support.

Last, but not least, I want to thank my family for their never-ending support and pride in my accomplishments. I am also deeply indebted to my husband Tim Wood for making the graduate school experience less lonely, less stressful, and more fun than it would have been otherwise. Thank you.

ABSTRACT

VARIATIONS ON STIGMERGIC COMMUNICATION TO IMPROVE ARTIFICIAL INTELLIGENCE AND BIOLOGICAL MODELING

SEPTEMBER 2011

MEGAN M. OLSEN

B.Sc., VIRGINIA POLYTECHNIC INSTITUTE AND STATE UNIVERSITY

M.Sc., UNIVERSITY OF MASSACHUSETTS AMHERST

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Hava T. Siegelmann

Stigmergy refers to indirect communication that was originally found in biological systems. It is used for self-organization by ants, bees, and flocks of birds, by allowing individuals to focus on local information. Through local communication among individuals, larger patterns are formed without centralized communication. This self-organization is just one type of system studied within complex systems. Systems of ants, bees, and flocks of birds are considered complex because they exhibit emergent behavior: the outcome is more than the sum of the individual parts. Emergent behavior can be found in many other systems as well. One example is the Internet, which is a series of computers organized in a self-organized fashion.

Complexity can also be defined through properties other than emergent behavior, such as existing on multiple scales. Many biological systems are multi-scale. For instance, cancer exists on many scales, including the sub-cellular and cellular levels. Many computing systems are also multi-scale, as there may be both individual and system-wide controls interacting together to determine the output. Many multi-agent systems would fall into this category, as would many large software systems.

In this dissertation I examine complex systems in artificial intelligence and biology: the growth of cancer, population dynamics, emotions, multi-agent fault tolerance, and real-time strategic AI for games. My goal is twofold: a) to develop novel computational models of complex biological systems, and b) to tackle key AI research questions by proposing new algorithms and techniques that are inspired by those complex biological systems. In all of these cases I design variations on stigmergic communication to accomplish the task at hand. My contributions are a new agent-based cancer growth model, a proposed use of location communication for removing cancer, improved multi-agent fault tolerance through localized messaging, a new approach to modeling predator-prey dynamics using computational emotions, and improved strategic game AI through computational emotions.

TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENTS	iv
ABSTRACT	vi
LIST OF TABLES	xii
LIST OF FIGURES	xiv
CHAPTER	
1. INTRODUCTION	1
1.1 Complex Systems	1
1.1.1 What are Complex Systems?	1
1.1.2 Complex Systems and Computer Science	2
1.2 My Approach	4
1.2.1 Modeling of Biological Systems	4
1.2.2 Inspiration from Biology to Computer Science	5
1.3 My Interdisciplinary Domains	7
1.3.1 Cellular Biology: Cancer	7
1.3.2 Cognition: Emotion	9
1.3.3 Population Dynamics	10
1.4 My Contributions	11
1.4.1 Mathematical Population Dynamics Model of Cancer	11
1.4.2 Agent-based Computational Cancer Model	12
1.4.3 Cancer-inspired Multi-Agent Fault Tolerance	13
1.4.4 Utilizing Emotions for Population Dynamics Modeling in Cellular Automata	14
1.4.5 Computational Model of Emotions for Real-time Game Artificial Intelligence	15
1.4.6 Communication Techniques	16

2. MAIA: MATHEMATICAL ANALYSIS OF CANCER THROUGH POPULATION DYNAMICS	19
2.1 Introduction	19
2.2 Related Work	20
2.3 Model	21
2.4 Dynamical Analysis	23
2.4.1 Minimum Death Value	23
2.4.2 Prediction Equations	25
2.5 Conclusions	30
3. AGENT-BASED CANCER AND COMMUNICATION MODEL	32
3.1 Introduction	32
3.2 Related Work	34
3.2.1 Spatial Cancer Models	34
3.2.2 Angiogenesis Models	35
3.3 Cancer Model	37
3.3.1 Genes and the Life Protocols	38
3.3.2 Nutrients	39
3.3.3 Angiogenesis	41
3.3.4 Model Validation	41
3.4 Communication Model	42
3.5 Simulation Details	46
3.6 Results	47
3.6.1 Communication Protocol Never Fails	48
3.6.2 Delay of Communication Protocol Start	49
3.6.3 Failure to Receive Messages	50
3.6.4 Failure to Send I'M DYING	51
3.6.5 Failure to Receive Messages and Send I'M DYING	52
3.6.6 Result Summary	52
3.6.6.1 Hypothesis 2.1	53
3.6.6.2 Hypothesis 2.2	53
3.6.6.3 Hypothesis 2.3	53
3.6.6.4 Hypothesis 2.4	54
3.6.6.5 Hypothesis 2.5	54
3.7 Conclusions	54
4. MULTI-AGENT FAULT TOLERANCE INSPIRED BY CANCER	57
4.1 Introduction	57

4.2	Related Work	59
4.3	HADES - The System	62
4.3.1	Agent Cloning	63
4.3.2	Internal Agent Repair	63
4.3.3	Maintaining Enough Agents	64
4.3.4	Presence Signals	64
4.4	Irregular Agents and the Communication Protocol	65
4.4.1	Rescue Protocol Communication	66
4.4.2	Double Messaging	67
4.4.3	Signal Parameters	68
4.5	Model Flow	69
4.6	Results	70
4.6.1	Simulation Details	70
4.6.2	Removal of Irregular Agents	71
4.6.3	Robustness	73
4.7	Conclusions	74
5.	EMOTIONS FOR PREDATOR PREY	77
5.1	Introduction	77
5.2	Previous Work	79
5.3	Our Model	82
5.3.1	Probabilistic and Neighbor-based Rules	84
5.3.2	Individual Emotions	85
5.3.3	Emotion Communication - Direct	87
5.3.4	Emotion Communication - Stigmergic	88
5.3.5	Rules Enhanced by Emotions	88
5.4	Results - Role of Emotion in Predator and Prey Decisions	90
5.4.1	Experimental Design	90
5.4.2	Population Dynamics	91
5.4.3	Rabbits	92
5.4.4	Foxes	93
5.4.5	Discussion	95
5.5	Results - Comparison of Communication Paradigms	96
5.5.1	Experimental Design	96
5.5.2	Effect on Emotion	97
5.5.3	Effect on Population	99
5.5.4	Discussion	100

5.6	Conclusions	101
6.	EMOTIONS FOR AGENT COORDINATION	104
6.1	Introduction	104
6.2	Related Work	105
6.3	The System	107
6.3.1	Globulation	107
6.3.1.1	Agent Types	107
6.3.2	Emotions	108
6.3.2.1	Types of Emotions Modeled	108
6.3.2.2	Emotion Sharing Map	109
6.3.2.3	Agents using Emotions	111
6.4	Experimental Design	113
6.5	Results	114
6.6	Conclusions	117
7.	CONCLUSIONS AND FUTURE DIRECTIONS	119
7.1	Conclusions	119
7.2	Future Directions	121
7.2.1	Cancer	121
7.2.2	Multi-agent Fault Tolerance	121
7.2.3	Computational Emotions for Predator Prey Modeling	122
7.2.4	Computational Emotions for Real-time Games	122
	BIBLIOGRAPHY	123

LIST OF TABLES

Table	Page
2.1 The four equations representing different models of global spatial requirements.	23
2.2 Approximation of minimal removal values for all equations based on results from simulation of the ODEs.	25
2.3 Fixed Points and their eigenvalues for the general form of the equations.	26
2.4 Fixed Point fx_1 ($C = 0$, $M = 0$) for each equation type, written in our original variables	28
2.5 Fixed Point fx_2 for each equation type, written in our original variables. This fixed point is only valid for these two cases.	28
2.6 Examples of fixed point values for different parameter values for the fixed point fx_3 when $k = 1$ and $N = 1000$. These values are obtained via simulation of the equations.	30
3.1 Cellular basic protocol parameter values used in this study. Each line will be referred to in the text by the ratio, which is the cancer proliferation rate divided by the healthy proliferation rate. Each cell may vary by up to 10% from its creator, but must remain within the given ranges.	46
3.2 Apoptosis threshold tested with each parameter set.	47
3.3 Parameter sets for the communication protocol.	47
4.1 Agent parameter values used in this study. Each line will be referred to in the text by the ratio, which is the irregular cloning rate divided by the healthy cloning rate. Each agent may vary by up to 10% from its creator, but must remain within the given ranges.	70
4.2 Parameter sets for the communication protocol.	70
4.3 TOS tested with each parameter set.	71
5.1 Experience values affecting each emotion. For each emotion the corresponding value of the experience variable is listed.	86

5.2	Non-emotion parameters for each species. Rep is the probability of reproducing after reaching the Maturity Age, Cure is the probability of being cured from, DiseaseMove is the probability of an individual moving if it is diseased, and Starvation is the hunger level that causes death.	91
5.3	Emotion parameters used in experiments. Decay rate linearly decrements the emotion value at every point, Rat denotes the ratio of how disgust and anger affect reproduction rate, $c_{c,sp}$ discounts communicated emotion from surrounding conspecifics, and the fear threshold denotes the level of fear necessary to pause rabbit reproduction.	91
6.1	Diffusion radii tested for fear and frustration.	114

LIST OF FIGURES

Figure	Page
1.1 Figure from Kitano, 2002 [79]. The ideal cycle between biology and computation. Biologists do their standard “wet” experiments, which creates new biological knowledge. Computational scientists then use that knowledge to inform their model, which through simulation creates new predictions that can feed into new standard biological experiments, and the cycle continues.	4
1.2 American Cancer Society 2010 cancer statistics figures on the state of cancer as a disease in the US [http://www.cancer.org/Research/CancerFactsFigures/CancerFactsFigures/cancer-facts-and-figures-2010]. Sources for each set of data is listed at the bottom of the figure.	7
1.3 copyright Hanahan and Weinberg, 2000 [65]. This figure depicts the main mechanisms that need to be in place for a cell to become a cancer cell. Many models are based on this definition and abstraction.	9
1.4 Variation on stigmergic communication. An ‘X’ represents the communicating entity, and a circle represents some of the locations in which neighboring entities will receive the communication. A dark background is a strong message and a light background is a weak message; white represents no shared message. The agent-based cancer model and the multi-agent fault tolerance message passing protocol both use the communication in (a), where messages are sent to the nearest neighborhood first and then a weaker message follows to the further neighborhood. Emotions for real-time strategic games are communicated with an immediate diffusion at every time step as seen in (b). The predator-prey model uses two different versions of communication. In (c) it does not diffuse, but instead entities at each time step check the emotions their immediate neighbors had the previous time step. In (d) the map remembers the emotion from the previous time step, although it is not diffused further from the cell where it originates.	17

2.1	Demonstration with Σ_3 and Σ_4 for $\rho_N = 0.01$, $\rho_C = 0.1$, $\mu = 0.00000001$, $\delta_N = 0.001$, $Z = 1000$. Solid lines represent the number of PopN cells, dashed lines represent the number of PopC cells. Changes in slope represent proliferation and death changes based on Γ . Each equation set ends in an equilibrium with PopN as the majority only when δ_C meets the minimum requirements for the equation set. For Σ_3 : (a) δ_C is too low, giving a PopC majority (b) δ_C is high enough to give a PopN majority. For Σ_4 : (c) δ_C is too low, giving a PopC majority (d) δ_C is high enough to give a PopN majority. Note that for Σ_4 PopC cells are constrained to the size of the system, whereas in Σ_3 they can increase infinitely.	24
2.2	Phase planes for the four equations and different parameter values, with red circles representing fixed points. (a) $\rho_N = 0.1$, $\rho_C = 0.1$, $\delta_N = 0.01$, $\delta_C = 0.01$, $\mu = 0.001$ (b) $\rho_N = 0.1$, $\rho_C = 0.1$, $\delta_N = 0.1$, $\delta_C = 0.01$, $\mu = 0.01$ (c) $\rho_N = 0.1$, $\rho_C = 0.1$, $\delta_N = 0.01$, $\delta_C = 0.11$, $\mu = 0.01$ (d) $\rho_N = 0.3$, $\rho_C = 0.3$, $\delta_N = 0.001$, $\delta_C = 0.01$, $\mu = 0.001$	27
3.1	Flow chart for both normal and cancer cells, for every time step. The differences for normal and cancer cells are what probabilities are used and how the genes affect decisions (not shown).	38
3.2	Cells cycle through the normal and hypoxic states, depending on the amount of nutrients they receive. Once they are not receiving nutrients they become necrotic and die.	40
3.3	Change over time in number of cancer (red) and healthy tissue (blue) cells. (a) When angiogenesis is not activated, cancer cells grow quickly, but are unable to grow large enough to crowd out healthy cells. (b) When angiogenesis is activated, the cancer cells are able to grow much larger and take nutrients and space away from healthy cells such that they basically disappear.	42
3.4	Change over time with (blue) and without (red) angiogenesis: (a) Number of tumor cells, (b) Number of hypoxic (low nutrient) cells, (c) Number of new vessel cells, and (d) Number of non-hypoxic tumor cells. More tumor cells are able to grow with angiogenesis, and eventually there is also more rapid creation of hypoxic cells with angiogenesis. These graphs are very similar to those in [138].	43
3.5	The PLEASE DIE signal is first sent due to a push occurring, as seen in the middle image of (a). Once a cell is ready to undergo apoptosis from these signals, it will send the I'M DYING signal as seen in the middle image of (b), and then will undergo apoptosis as seen in the last image of (b).	44
3.6	Flow chart for both normal and cancer cells including communication of I'M DYING and PLEASE DIE signals, for every time step. The differences for normal and cancer cells are what probabilities are used and how the genes affect decisions (not shown). Receipt of I'M DYING and PLEASE DIE signals is not shown as that happens in a separate process between ticks.	45

3.7	Rate of success in removing cancer cells. All proliferation ratios and apoptosis thresholds generate the same results.	48
3.8	Rate of success in removing cancer cells when there is a delay in the start of the communication protocol. The legend shows the parameter set. All apoptosis thresholds are combined, because there is no variation among them.	49
3.9	For each proliferation ratio, the highest rate of failure to receive messages that still results in 100 percent removal of cancer cells (diamond), and the lowest rate of failure to receive messages that results in no removal of cancer cells (circle). For all ratios, the fourth (strongest) parameter set can succeed when tumors do not receive messages up to about 77% of the time.	50
3.10	For each proliferation ratio, the highest rate of failure to send messages that still results in 100 percent removal of cancer cells (diamond), and the lowest rate of failure to send messages that results in no removal of cancer cells (circle). For all ratios, the fourth (strongest) parameter set can succeed when tumors do not send I'M DYING up to about 90% of the time.	51
3.11	For each proliferation ratio, the highest rate of both failures that still results in 100 percent removal of cancer cells (diamond), and the lowest rate of both failures that results in no removal of cancer cells (circle). For all ratios, the fourth (strongest) parameter set can succeed when tumors fail to adhere to the messaging rules around 60% of the time.	52
4.1	The basic action controls for a healthy agent at each time step.	69
4.2	Percentage of experiments with success, defined by all irregular agents removed without all normally functioning agents removed. See Table 4.6.1 for parameter sets. Each bar represents a different threshold for dying from signaling, denoted by "normal/irregular."	72
4.3	The percent of experiments with success on the most successful dataset when irregular agents either ignore signals received (blue) or do not send I'M DYING messages (green). The system can still succeed in removing all irregular agents if either irregular agents ignore at most 68% of messages received, or they do not send I'M DYING messages 72% of the time.	73
4.4	The percent of experiments with success on the most successful dataset when signals are not used until some percentage of the system has been overrun with irregular agents. Even when waiting until 50% of the system has irregular agents, the system is still able to remove them completely.	74
5.1	Simulation world at multiple time steps shows the dynamics of model, and how it may differ between a world (a) without emotion, (b) without fox emotion but with rabbit emotion, (c) with fox emotion but without rabbit emotion, and (d) a world with fox and rabbit emotion. Vacant squares are in black, carrots are in dark gray, rabbits are in gray, and foxes are in white.	83

5.2	Population dynamics of the system with different parameter settings. In (a) all population sizes tend towards non-zero attractors. In (b) the fox population crashes while the rabbit and carrot population sizes continue to tend towards non-zero attractors.	84
5.3	Probabilistic and Neighbor-based Rules for rabbits and foxes.	85
5.4	The two communication styles. The 'x' represents where an entity currently exists, and the colored background represents the currently shared emotion at that cell (darkest is strongest). Both images represent the state of the shared emotion from that entity after three time steps.....	87
5.5	Some emotions modify fox and rabbit behavior negatively, and others modify the behavior positively. Both movement and reproduction are affected by emotion.	89
5.6	Population variation over time. The population over time is shown for no emotion, rabbit-only emotion, fox-only emotion, and both species emotion. The rabbit population (a) is highest when only rabbits have emotion with no emotion as a close second, and lowest when only foxes have emotion, and all emotion as second worst. The fox population (b) and carrot population (c) are both highest when only foxes use emotion, and lowest when only rabbits use emotion, with all emotion and no emotion close in the middle.	93
5.7	Rabbit consumption, disease, and starvation are all correlated. (a) Rabbits eat most frequently when foxes use emotion. This graph is the inverse of Figure 3(a), as rabbits eat the most when they have the lowest population size. (b) Rabbits are the most diseased with fox-only emotion, and the least diseased with rabbit-only emotion. Disease is highly correlated to amount of carrots eaten. (c) Although rabbits starve at a very low rate overall, they starve the fastest in the scenario where they eat the least and are the most diseased. These trends are very similar to the trends in Figure 5.6(a).	94
5.8	(a) Rabbits reproduce most frequently with fox-only emotion, closely followed by no emotion. Rabbits reproduce the least when they use emotion, likely due to fear, disgust, and anger significantly decreasing their probability for reproduction. (b) Foxes reproduce most frequently when they do not use emotion, most likely due to disgust and anger decreasing their reproduction probability. The effect is less pronounced than for rabbits as there is no fear level that stops them from reproducing completely.	95
5.9	Fox consumption, disease, and starvation are correlated. (a) Foxes eat the most frequently when they use emotion. (b) Foxes are the most diseased when they use emotion, and thus when they eat the most. Thus, emotions are not improving a fox's ability to avoid diseased food. (c) Foxes starve fastest when they do not use emotion, corresponding to when they eat the least and are the least diseased. Starvation rate does not directly correspond to population size.....	96

5.10	Average rabbit (a) and fox (b) populations for each parameter combination. The x-axis shows the communication coefficients (0.1,0.5,0.9) and decay values for each type of communication (N: no communication; D: direct; S: stigmergic (L=linear, G=geometric)). Black (darkest) represents when both species use emotion, Red (second darkest) represents when only foxes use emotion, and Blue (lightest) represents when only rabbits use emotion.	97
5.11	Average individual rabbit emotion for each parameter combination. The x-axis shows the communication coefficients (0.1,0.5,0.9) and decay values for each type of communication (N: no communication; D: direct; S: stigmergic (L=linear, G=geometric)). Black (darkest) represents when both species use emotion, Red (second darkest) represents when only foxes use emotion, and Blue (lightest) represents when only rabbits use emotion.	98
5.12	Average individual fox emotion for each parameter combination. The x-axis shows the communication coefficients (0.1,0.5,0.9) and decay values for each type of communication (N: no communication; D: direct; S: stigmergic (L=linear, G=geometric)). Black (darkest) represents when both species use emotion, Red (second darkest) represents when only foxes use emotion, and Blue (lightest) represents when only rabbits use emotion.	103
6.1	The plane representing the range of a unit's emotions and 4 possible emotional stages: the origin is no fear or frustration, representing contentment; point 1 represents a unit with little frustration but high fear; point 2 is a unit with low fear and medium frustration; and point 3 represents a unit with high frustration and high fear.	108
6.2	A series of images demonstrating a player's Emotion Sharing Map changing over time. Each of the 6 images represents the entire game environment. Images are taken every 4,000 time steps. Frustration is shown in yellow(middle shade of gray), Fear is shown in red (darker shade), the overlap of the two emotions is green (lightest shade), and the lack of shared emotion is black. Images are organized chronologically from left-to-right and top-to-bottom.	109
6.3	Approximate diffusion concept. The map in 6.3(a) depicts the values in the squares under and surrounding an agent that just experienced an event that resulted in a total emotion value of 10. If emotions decay linearly by 2, the map in 6.3(b) depicts the values in those same squares after a single time step before agents send their emotions to the map again.	111
6.4	The decision tree for a warrior at each time step.	112
6.5	Percent of the eight games won by NicowarESM (Player 1) against Nicowar (Player 2) with varying diffusion radii (as labeled). The dashed line shows the baseline, i.e. the percent of wins by Nicowar when playing itself without emotion.	115

6.6	Difference of average hp per agents and buildings from Nicowar (player 2) versus NicowarESM (player 1) with varying diffusion radii (as labeled). (a) When examining all experiments, we see that the most significant increase in HP is with diffusion radius of (2,2). (b) If we only examine cases where player 1 won the game, not only do most average HPs increase, but also the diffusion radius of (3,2) increases higher above Nicowar's HP.	116
-----	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----

CHAPTER 1

INTRODUCTION

1.1 Complex Systems

1.1.1 What are Complex Systems?

The term *Complex Systems* has been used in a variety of ways by scientists, often referring merely to a system that is not simple. However, the field of Complex Systems studies a much less general class of problems that are complex in ways that may include simplicity at its heart: emergent, adaptive, nonlinear, chaotic, and multiscale complexity, to name a few. These problems describe much of the world around us, and also apply to computing problems.

Emergent behavior refers to the whole as being more than just the sum of the parts, for example, when simple rules create complex dynamics. Stated another way, as the scale of the problem increases, so does the complexity, in a way that cannot be explained simply by the fact that there are more elements. Learning the fundamental laws (such as in elementary particle physics) does not necessarily provide the ability to reconstruct the world. Anderson summarizes this by stating that “the reductionist hypothesis does not by any means imply a ‘constructionist’ one” [8]. However, learning what creates this complexity within systems, as well as how these complex systems work and interact with each other, can further our understanding of many fields and of science in general.

Self-organization is often a noticed outcome of emergent behavior, and a prime example of complex behavior appearing from simple rules of interaction. It can be seen in many biological systems, including molecular or cellular biology, dynamics of species in ecology, or human behavior. A popular example is swarm intelligence, which refers to the natural organization of bees, ants, or flocks of birds without centralized communication [17, 30, 145]. In this case, self organization through simple interactions are causing a large coordinated effort for travel, finding food, or general communication [146]. The behavior emerges from the interactions of these simple rules, to create behavior more complex than would be expected by the rules themselves. This phenomenon is also observed in computer networks. Probably the best-known example is the Internet, a network that is

organized without central control (i.e. it is self-organized) but yet still follows a power law in link organization [112]. This structure of the network may affect more than just its organization, such as how it evolves over time and how data is shared through it [112, 99].

It has been shown that only a small amount of information from a few informed individuals is necessary for effective group decision making [32]. Communication in which only a relatively small amount of information is shared through the environment, which is then used to inform entity decisions, is called *Stigmergy*. Stigmergy is used by ants and bees as a way to enable self-organization. It can also be used for human collaboration, or to encourage collaboration and cooperation in computing systems. Wikipedia is an example of stigmergic communication by humans, where each individual is only modifying their own local information when they update a page on which they are an expert. However, all of those local updates combine to create a large functioning encyclopedia.

Systems that exhibit emergent behavior and self-organization may also interact on *multiple scales*, as do systems that we do not traditionally call emergent. Many systems in both biology and computer science work on multiple scales, including the scales of time and space. Cellular biology is a prime example of a multiscale system: even if we only look at what controls the cell, and ignore the overall interactions at a cellular level, there are different time scales for replication and interaction, and different sizes for proteins, RNA, etc. Thus, studying cellular biology can be done by examining cells at the cellular level, the subcellular level, or the tissue level. At the same time, it may be beneficial to analyze the overall functioning of the plant or animal created by these cells, or even the interaction of animals within a species or between species. Thus, every organism is representing multiple space/time scales of interest, and to truly understand the mechanisms at work it is necessary to analyze multiple scales at once. This leads to multi-scale models, both mathematical and computational.

1.1.2 Complex Systems and Computer Science

The field of complex systems is itself interdisciplinary by nature [82]. The goals of understanding complex phenomena in terms of organization, structure, and function are often served by both mathematical and computational models. Some of these models are similar to what computer scientists find familiar, such as differential equations, agent-based modeling, discrete-event modeling, or network analysis [5, 77]. Modeling provides a mechanism by which to represent phenomena and

attempt to understand what are the underlying mechanisms controlling that phenomena in nature [135].

Computer science has been captivated by many of these Complex Systems problems even if we do not generally think of them as being within the field of complex systems. One instance is the Game of Life defined by Conway, a set of simple cellular automata rules that leads to such complex behavior that the system is still not entirely understood [55]. Additionally, many have been investigating computational evolution through modeling as far back as Turing, who in his later years developed a mathematical model of the role of chemicals in biological development [151]. This branch of research continues primarily in the Artificial Life field, for instance using diffused chemicals to evolve a French flag [97] or for development from a single starting entity [126].

Artificial Intelligence has been moving in the direction of adaptive, or learning, systems for many years. One example is the field of reinforcement learning, which is inspired by behavioral psychology and used both to enhance computer systems and to model biological processes [147]. Complex adaptive systems are another subset of this direction, where we design new systems capable of learning from its experiences and able to adapt to different situations. These systems are often multi-scale, and can be used as a method of modeling complex systems to understand more than the basic rules leading to emergent behavior. Swarm intelligence is also used in the development of collaborating groups of robots and agents, both for self-organization and the ability to adapt to failures within a single agent [63, 28, 70].

Overall, computer scientists can offer much to the field of complex systems: modeling techniques, algorithmic development, insights from similarities between computing and non-computing problems, and machine learning, to name a few. Computer science itself is becoming more interdisciplinary as there is a move to work on interesting computational problems in other domains: linguistics, cellular biology, molecular biology, neuroscience, and social science are some of the most popular and emerging collaborations. From these collaborations, models are created and data is analyzed to aid in understanding the complexities within those fields, often arising from complex systems. What we learn from working with experts in these fields can also be used in the creation of new systems within computer science itself. Sometimes this refers to new algorithms or data analysis techniques, but it can also refer to new systems that are more distant from the original interdisciplinary interaction. In this dissertation I will investigate how stigmergic communication can




Figure 1.1. Figure from Kitano, 2002 [79]. The ideal cycle between biology and computation. Biologists do their standard “wet” experiments, which creates new biological knowledge. Computational scientists then use that knowledge to inform their model, which through simulation creates new predictions that can feed into new standard biological experiments, and the cycle continues.

increase collaboration between artificial entities, both when modeling non-computing systems and when building computing systems.

1.2 My Approach

My goal in this dissertation is twofold: a) to develop novel computational models of complex biological systems, and b) to tackle key AI research questions by proposing new algorithms and techniques that are inspired by complex biological systems. I examine two complex systems from outside of the computing domain from both the modeling and inspiration perspective: cancer and emotions. In some of this work I also incorporate the complex systems field of population dynamics as part of my approach. I investigate how communication can play a crucial role in more accurate modeling, and in development of an inspired system.

1.2.1 Modeling of Biological Systems

Modeling of biological systems allows us to analyze and approach them in a different way than is traditionally done by biologists. Ideally, biologists and modelers work together to form a complete cycle where biological experiments feed into computational/mathematical models, which then feed

back into the biological experiments to continue the cycle (Figure 1.1). Both fields are still working to achieve this collaboration and balance. However, many modelers are either also involved in the traditional “wet” biological experiments or work with biologists and their data to ensure that the created models are biologically validated, and can therefore have the potential to feed back into the biological experimental cycle. Ideally, models are used for their predictive power such as is currently attempted in epidemiology for predicting the spread of disease with and without vaccines [155], and in geosciences for the prediction of earthquake activity [133].

There are many different types of models used to study biological problems, which are usually complex systems: ordinary differential equations, partial differential equations, agent-based simulations, event-based simulations, cellular automata, computational geometry, and more. Non-linear dynamic approaches examining self-organizing systems are also used for multi-scale problems such as cancer [29]. The key is to determine the appropriate approach for the subfield and specific problem being studied.

In this dissertation I examine models of two biological problems: the growth and removal of cancer cells, and predator-prey dynamics in an ecological setting. I examine the role of space and healthy tissue cells in cancer growth using ordinary differential equations. I also examine the role of communication between those cells in the removal of cancer cells via agent-based modeling. For predator-prey dynamics I investigate the use of computational emotions and conspecific (intra-species) communication to model the interactions of entities within a species and how that effects the dynamics between competing species.

1.2.2 Inspiration from Biology to Computer Science

In addition to the traditional approach of modeling complex systems, I also take this newly gained knowledge and use it as inspiration in improving computing systems. In this case, the goal is to create a biologically inspired algorithm or system that solves a problem within artificial intelligence. To create this inspired algorithm or system, it is necessary to draw parallels between the biological system and some computing system. Then, a problem within that computing system should appear related to a mechanism within the biological system. At this point, a new algorithm for the computing system can be determined based on what is known about the biological system. One of the classic examples of biologically-inspired algorithms is genetic algorithms [60]. This

class of algorithms allows a system to learn the best state through continuous mutations and fitness analysis. Genetic algorithms are based on the idea of biological mutations that lead to natural selection over time.

Another example of biologically-inspired computing is wireless sensor networks designed with self-healing capabilities inspired by immunology to detect sensor faults. SASHA mimics B-cells in the immune system with scripts on monitor nodes that follow the status of the sensor nodes in the system [16]. The monitor nodes can find failure by examining the statistical properties of the sensor readings. This system can adapt to the changes in the network caused by sensor failure via monitor nodes notifying sensors of incorrect readings, allowing them to request retraining. This combination of different node types interacting enables the system to find and react to failures [16].

Hardware can also be developed through biological-inspiration, although it is less common. In [88], regeneration, repair, and death are combined to create an artificial organism as the first step toward hardware with the ability to remove surrounding agents that may be faulty. The cells in the system are organized in a grid. One form of repair involves disabling all cells in the column of the faulty cell after transferring their functions to the cells in the column to their right, essentially using death to repair the organism. The other form of repair is internal, used to combat failure of the artificial molecules that control cellular actions. This repair is accomplished by removing the faulty molecule and then rearranging the remaining ones until a spare is reached, ending with the same number of molecules as was used before the failure [88].

Machine learning techniques have also been developed based on swarm intelligence [66]. In this case, data is clustered and sorted in an emergent behavior fashion that relates to the interactions of ants and flocks of birds. However, this work is not yet viable within machine learning but is primarily interesting for the novelty of the biological inspiration [66].

I approach biologically-inspired computing from the software perspective. In this dissertation I discuss two different biological inspirations: a cancer inspired multi-agent system fault tolerance communication protocol, and computational emotions for real-time strategic game artificial intelligence.

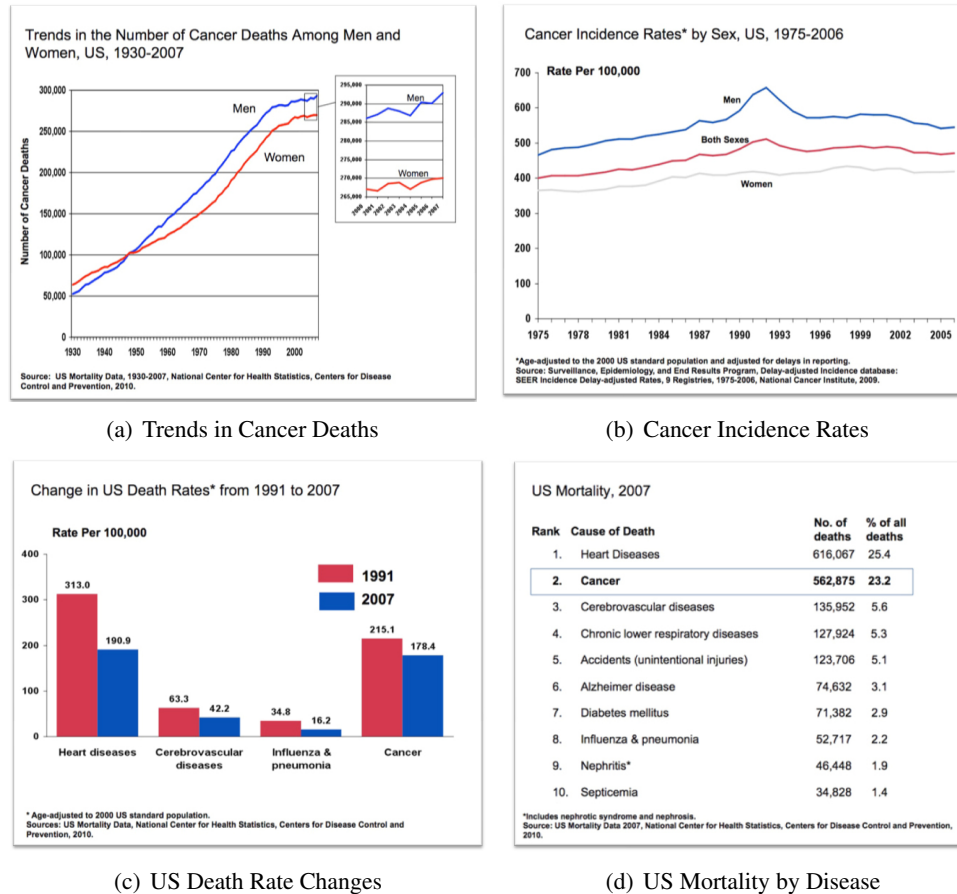


Figure 1.2. American Cancer Society 2010 cancer statistics figures on the state of cancer as a disease in the US [<http://www.cancer.org/Research/CancerFactsFigures/CancerFactsFigures/cancer-facts-and-figures-2010>]. Sources for each set of data is listed at the bottom of the figure.

1.3 My Interdisciplinary Domains

1.3.1 Cellular Biology: Cancer

Cancer incidence is expected to rise worldwide from 12 million new people affected annually in the year 2000 to an anticipated 20 million in the year 2030, highlighting the urgent need to identify highly effective preventative and therapeutic interventions. Not only are more people being diagnosed each year (even when adjusted to a standard population size), but there are also more total deaths from cancer. These statistics from the American Cancer Society can be seen in Figure 1.2. Additionally, there is less progress being made overall in fighting cancer, despite the amount of press that the disease has (Figure 1.2(c)). The only disease that is currently more deadly than cancer

is heart disease; but in five years cancer is expected to be top of the list for the first time (Figure 1.2(d)).

The overall goal of cancer research is to either eradicate the disease or make it no longer deadly. This goal is broken into many smaller high-level goals such as understanding how cancer cells are formed, how a single cancer cell develops into a tumor, what regulates the growth of cancer cells, how different types of cancer differ, how to better test for cancer in a patient, and how to remove cancer once it is discovered. The fields of Computational Biology and Systems Biology work toward a better understanding of the mechanisms behind this disease in a way that standard biological techniques cannot yet achieve. With the use of computational and mathematical models we can examine how biological systems work in combination, as well as create new testable hypotheses [154]. These models must be based on biological knowledge and data and validated before their results can be used to direct biological research. Although there is great promise from computational and mathematical techniques for cancer research, the cancer biology community is still in the process of learning to accept the potential of these models to contribute to the field [153].

To understand how to model cancer it is first necessary to have a better understanding of the disease itself. Cancer is essentially a disease in which certain cells no longer follow the originally defined genetic rules. There are a number of aspects to biological cells that are important for sustaining a cellular system such as a tissue. Cells are not only controlling their own behavior, but through inter-cellular communication they are also controlling the behaviors of their neighbors. Additionally, there are many larger aspects within the body that affect the cells, such as nutrient diffusion. Hanahan and Weinberg defined a set of six properties that represent the failures needed within a cell for it to become cancerous [65]. All of these modifications to the cell and tissue are necessary for a metastatic tumor cluster to form (Figure 1.3.1).

However, the order in which these failures occur is also important for a cell to become cancerous. These failures occur via mutations to the cell's DNA, which can occur due to a variety of causes including during proliferation, i.e. the creation of a daughter cell. Thus, the ability to repair genetic mutations and the ability to undergo apoptosis if that repair fails must be some of the first mutations to occur for a cell to become cancerous. Afterward, the cell must become mutated in the tumor suppressor genes (p53), as otherwise the detrimental effects of increased growth will not occur. Next, the cell may be mutated in the ability to proliferate such that its proliferation is increased.

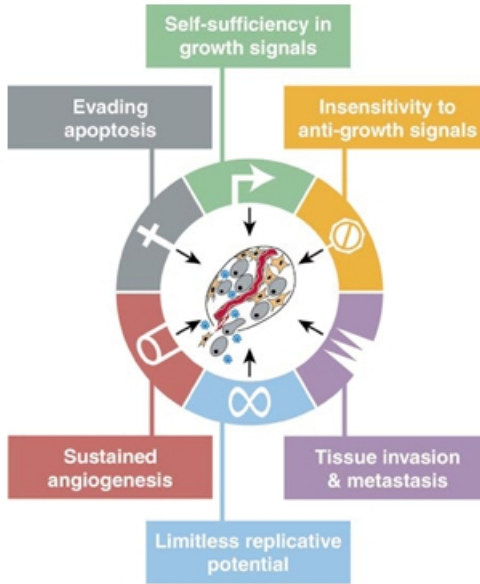


Figure 1.3. copyright Hanahan and Weinberg, 2000 [65]. This figure depicts the main mechanisms that need to be in place for a cell to become a cancer cell. Many models are based on this definition and abstraction.

This may include multiple mutations as the rate of proliferation is controlled in multiple ways. Mutations in a different order will not create a cancer cell, but instead will create a cell that will undergo either repair or apoptosis (self-death) to maintain the system [117, 51, 85]. An overview of the genetics of cancer progression is provided by Michor [96].

With these mutations the cell has accomplished four of Hanahan and Weinberg's requirements. Additionally, for a tumor to grow large enough to be problematic it must be able to induce angiogenesis, or the creation of new blood vessels. These new blood vessels will bring new nutrients to the cancer cells, allowing the tumor cluster to grow larger. Eventually these vessels may allow the tumor cells to travel to other parts of the body to create new tumors elsewhere. This process is known as metastasis, and is what makes many cancers both dangerous and difficult to remove.

1.3.2 Cognition: Emotion

Emotions are one of the many aspects of the human psyche that are still in the process of being understood. Emotions are studied both in psychology, as well as computationally. The goals in emotion research vary between locating the part of the brain responsible for emotions, determining how emotions affect us, and creating computational representations of emotions.

Rolls views emotions as states that are elicited by rewards and punishments, and that any reward or punishment will elicit an emotion [124]. Similarly, he theorizes that emotions can be caused not only by external stimuli, but also from recalling emotional events. These emotions also perform functional roles when they occur: eliciting autonomic or endocrine responses such as increased heart beat or adrenaline, motivation, communication, social bonding, episodic memory (strong emotions will increase storage, but may not guarantee accuracy), and memory recall [124].

There is much debate on exactly what constitutes an emotion and how that is different from a mood [124]. There is also discussion on what constitutes a basic emotion (innate) versus a complex emotion. Many lists of the basic emotions exist, varying from the three emotions of pleasure, pain, and desire from Spinoza [143, 35] to the original six basic universal emotions from Ekman of happiness, sadness, fear, anger, surprise, and disgust, and then to Ekman's more recent list of seventeen basic emotions [45].

Even after we agree on the definition of basic or complex emotions, we must then determine how to classify them. Russell suggested that all emotions have both valence (positive versus negative) and arousal (degree of intensity) [127]. Additionally, Plutchik states that the basic and complex emotions can be visualized on a three dimensional cone, where emotions are organized by intensity and similarity in eight dimensions [118]. His basic emotions form the middle level of intensity, with one set of complex emotions being more intense, and another set being less intense.

In the last few years it has been suggested that emotions constitute an important part of adaptive decision making systems, contradicting the older view that emotions typically interfere with decision making [36, 132]. Case studies reported that people who suffered injury to or loss of areas of the brain related to emotion also experienced impaired decision making [13]. It thus seems likely that emotions are a key factor in human interactions and decisions.

1.3.3 Population Dynamics

Population Dynamics is interdisciplinary by definition, as it is a mathematical field used to study the development of both a single as well as multiple interacting species. The techniques have been applied in various biological fields, most notably in ecology, epidemiology, and cellular biology. In ecology, models are used to show the change in plant and animal populations, such as which trees

will survive in a forest over many hundreds of years, or what ratio of species is sustainable. Many techniques in population dynamics originate from studying ecological systems.

Epidemiology models are used to describe the spread of diseases. Although there are specific epidemiological models such as SIR, the techniques used in ecology can be similarly applied in epidemiology. Likewise, when studying how cancer cells may grow in a system, population dynamic approaches can be utilized to determine how quickly cells grow and how they may interact with cells within their neighborhood. Although these application fields consider different populations, many techniques can be shared among them for modeling the interaction and development of the populations.

1.4 My Contributions

This dissertation examines both modeling specific problems within the interdisciplinary domains, and developing computational tools that can be further applied to other problems within the same domain. For each domain I examine how local communication can play a role in either the biological system or the computational system developed to examine the biological one. In this section I discuss five claims specific to individual domains, and one claim that is general across all domains, which will be supported by this dissertation.

1.4.1 Mathematical Population Dynamics Model of Cancer

I have already described how cancer is a complex problem studied on many scales. One way to study cancer is to examine the healthy and tumor cells as competing populations. In this case there are multiple resources being competed for, primarily nutrition and space. These two resources are inherently intertwined, as too many cells in an area will both provide too little space and too few nutrients to neighboring cells. I thus propose modeling the growth of cancer as a population dynamics mathematical model to determine how simple rules may be able to capture the fundamental aspects of cancer growth.

Claim 1 *Varying spatial regulations within a cellular system leads to significantly different number of predicted cells in a model of cancer cell growth.*

The necessity of including spatial requirements among entities such as with the growth of trees or the spread of disease through a population, as well as how to adequately represent space, has been a

topic of debate in population dynamics modeling [83, 75]. Some scientists believe that providing the exact coordinates of the involved entities is crucial for accurate results, whereas others have argued that an approximation of space, such as global density, is sufficient. Although spatial requirements have not been traditionally examined as a mechanism affecting the competition between healthy and cancer cells in a tissue, I believe it may play a crucial role. I therefore analyze how spatial availability affects the growth of cancer cells in an ODE system in which I also consider the growth of healthy cells [108].

1.4.2 Agent-based Computational Cancer Model

Agent-based models are ideal for modeling cancer at the cellular level because they can provide a spatial environment, a specific neighborhood, and cell centric controls. I develop an agent-based cancer model that includes both healthy and tumor cells. Only recently have researchers begun analyzing models in which healthy and cancer cells interact with each other explicitly. Often, healthy cells are either ignored or factored in implicitly. However, modeling those cells in addition to cancer cells creates a more robust and accurate picture of the factors affecting tumor growth. Cells are able to move within their three-dimensional environment, proliferate, and undergo apoptosis. Cells make decisions based on probabilities and their environment, including whether they have enough nutrients. This model is validated by comparison to published results.

Claim 2 *Intercellular messaging among cells based on neighbor death and spatial impingement can be used to encourage death of surrounding cells such that primarily cancer cells are killed and healthy tissue cells survive.*

I use this model to then test if inter-cellular communication of “death” signals that induce apoptosis will allow the tissue to fight back against the tumor cells. These signals are based on signals that have been found in biological experiments. I test how and when this type of signal may be used successfully against cancer cells, and suggest that they be further studied by biologists to complete the cycle of collaboration. If these signals work as predicted by the model, then there is the potential for them to eventually be used in therapy against cancer [109].

1.4.3 Cancer-inspired Multi-Agent Fault Tolerance

I investigate fault tolerance for a system of cooperating agents. For a multi-agent system to function continuously it must adapt on-line to failures. There are essentially three different ways in which a system can fail: unreliable infrastructure, non-compliant agents, and emergent dysfunctions. Although each of these types of failures may cause different problems for the system, in each case the same high level process needs to be followed to deal with the problem: need to acknowledge and diagnose the problem, and then fix the problem. I will focus on the approaches to fixing the problem.

Overall there are two main approaches for fixing these three types of problems in multi-agent systems: Survivalist and Citizen. The survivalist approach requires each agent to be capable of dealing with all problems as an individual following a prepared set of actions for each specific problem [90]. For instance, a set of replicas could be maintained and then deployed once a fault is detected. However, this approach requires the designer to be able to anticipate all types of faults within the system and cannot easily deal with agent death. The citizen approach is the other extreme, as it utilizes an external system that is alerted when an agent dies and then reallocates tasks so that the overall system continues to function correctly [80]. Thus, the Citizen approach can very easily deal with agent death. However, there is now an additional system that must be maintained, creating another point of failure.

My approach is a combination of these two techniques as a Citizen Group approach. It differs from the survivalist approach as it does not require all agents to deal with failures individually. Unlike the citizen approach that requires special monitor nodes to diagnose failures, my system has each agent monitor its neighbors to detect and eliminate anomalies. This citizen group approach is accomplished by defining a cancer-inspired mechanism for multi-agent systems that improves robustness by enabling agents to combat anonymous malfunctioning agents.

Claim 3 *The intercellular messaging investigated for cancer removal can improve multi-agent system fault tolerance by allowing agents to use only local information and collaboration to remove faulty agents.*

The cancer model shows how local communication can be utilized to remove cancer cells without necessarily knowing which cells were cancerous. However, based on the fact that cancer cells grow in a cluster this mechanism removes all cancer agents by only detecting the irregularities of a

few of them. Seeing that this mechanism works well to remove cancer, I apply the same techniques to the problem of multi-agent fault tolerance. In this case, I examine agents instead of cells, and determine how to remove malfunctioning agents instead of removing cancer cells [110, 107].

1.4.4 Utilizing Emotions for Population Dynamics Modeling in Cellular Automata

I introduce intra-species disease transmission and emotion-inspired rules for our predator and prey (foxes and rabbits). Real populations in nature are subject to epidemic diseases, a number of which can cross species. Such diseases have significant effects at the level of individual behavior and population dynamics. Evidence suggests that a primary contributor to the evolution of the emotion disgust is protection from the risk of disease [34]. I explore the relationship between disease transmission and emotional response. For collective behavior to arise information is shared between conspecifics and individual decisions are made on that information. The information shared is in the form of emotions, and both rabbits and foxes make decisions with their emotions and the shared emotions taken into account. Emotions are affected by environmental events, and thus represent a high level of information about the environment. The development of emotions in higher animals has been conjectured to originate for purposes of survival in basic scenarios such as predator-prey [15, 87], and thus emotionally-inspired rules are a natural extension to the traditional CA framework. Although they have been suggested previously for CA [3], I am unaware of any work utilizing emotions in the context of predator-prey dynamics modeled within a CA framework.

Claim 4 *Computational emotions provide a framework for modeling predator-prey dynamics that will provide different modeling behavior than a traditional cellular automata model.*

Thus, I choose to include the six fundamental emotions as defined by Ekman [45] to our rabbits and foxes, namely happiness, sadness, fear, anger, disgust, and surprise. Emotions occur in response to specific world events, such as the happiness of food consumption and the fear of predator encroachment. Additionally, I enable conspecific communication of emotions to aid in coordination and cooperation. In other words, the emotional state of a member of a species will be communicated to a member of the same species within a restricted surrounding, and affect their emotional state. I analyze two approaches to communication: a direct communication where information does not

linger past one step in time, and a slowly decaying trail similar to ant pheromones. In both cases entities only utilize local information.

I consider this approach to emotional communication as an efficient way of transferring information that is crucial for the survival of the group. Our analysis shows that emotions improve the population sizes for both predator and prey. I show that it is in the best interest of both rabbits and foxes to use emotion if they do not know if the other species will use emotion, and that the choice of decay function changes the behavior of the populations [105, 106].

1.4.5 Computational Model of Emotions for Real-time Game Artificial Intelligence

Real-Time Artificial Intelligence (AI) has been investigated for over a decade [100]. A system is considered to be a Real-time AI system if it is able to make decisions within a guaranteed response time and thus meet domain deadlines. These systems face many challenges, including working with partial information, choosing the most crucial action if there are multiple scenarios to react to, and working continuously for an extended period of time without failure. These systems are usually created as expert systems, as they are used for a specific domain.

Real-Time Strategy (RTS) is an offshoot of general purpose real-time AI. RTS refers specifically to systems where the primary purpose is to create strategy, usually in a competitive atmosphere. Although they may at first seem unrelated, emotions can play a large part in strategy especially when time is limited. Emotions are believed to improve our response time, increase our memory capacity, and provide quick communication [124]. We are able to notice things that we fear quicker than things we enjoy or are indifferent about, showing fear to be crucial to our response time. Emotions help us convey our experience to another person; for instance, they will realize danger quicker from noticing our fear than by hearing our explanation.

Claim 5 *Computational emotions can be used to improve the performance of a computer player in a real-time strategy game.*

I propose including emotion with RTS algorithms to enhance our strategy. Our system utilizes a current RTS gaming engine called Globulation that includes computer controlled agents. The computer tells these agents where to move, what to do, and when to create more of them; the same actions controlled by a human player. I provide computational emotions for these agents, and determine how those emotions affect the game play. One of our main contributions is the creation

of an Emotion Map that enables units to communicate their emotions in a way similar to Stigmergy. This Emotion Map saves the emotion of units and diffuses it for a period of time, enabling other units to feel the emotion of their peers. Thus, there is an indirect communication between a single player's agents. If the emotions are designed to be reactive to the environment, this map would enable agents to lay a trail for moving to or from specific types of areas without the need for either detailed or direct communication [104].

1.4.6 Communication Techniques

Claim 6 *Stigmergic communication can be utilized to improve collaboration in both modeling interdisciplinary problems and when designing computational systems.*

In each of the above contributions I examine how different forms of communication based on stigmergy can increase collaboration between entities. In the agent-based model of cancer I develop two diffused message systems: one for signals that are constantly sent by cells, and another for signals that are only sent by cells when a specific event occurs. This latter type of communication can be seen in Figure 1.4(a). Although the messages are diffused over time within a specific radius, they do not linger past the time in which they are sent. Thus, only the cell currently in that location when the message is diffused there will receive the message. I examine how variations in the diffusion radius, signal strength, and interpretation of the second type of signals modifies cellular behavior. Both of these signals are then applied to multi-agent fault tolerance.

Two other variations on communication are compared for sharing emotions to improve predator-prey cellular automata modeling: direct communication that is constantly sent to neighbors within a radius of one (Figure 1.4(c)), and a temporally decaying trail mimicking ant pheromones (Figure 1.4(d)). Direct communication is accomplished by each entity sharing their emotion on their own location. Each entity at the beginning of each time step uses the information in their neighboring squares to affect their own emotion as well as their movement. The shared emotion then disappears before the next time step. The decaying trail is created by each entity sharing their emotion at their location without it immediately disappearing, but decaying over time. This communication is closer to stigmergic communication. I examine how this variation in communication paradigm affects the population size of each species within the system.

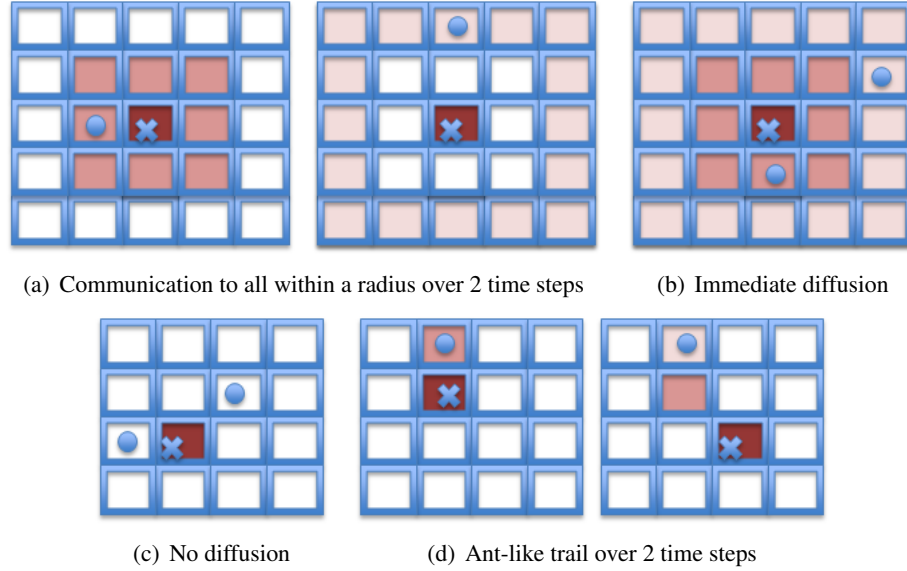


Figure 1.4. Variation on stigmergic communication. An 'X' represents the communicating entity, and a circle represents some of the locations in which neighboring entities will receive the communication. A dark background is a strong message and a light background is a weak message; white represents no shared message. The agent-based cancer model and the multi-agent fault tolerance message passing protocol both use the communication in (a), where messages are sent to the nearest neighborhood first and then a weaker message follows to the further neighborhood. Emotions for real-time strategic games are communicated with an immediate diffusion at every time step as seen in (b). The predator-prey model uses two different versions of communication. In (c) it does not diffuse, but instead entities at each time step check the emotions their immediate neighbors had the previous time step. In (d) the map remembers the emotion from the previous time step, although it is not diffused further from the cell where it originates.

Finally, I develop a hybrid of these communication techniques for sharing emotions in a strategic real-time strategy game. This last version of communication is constantly sent and temporally decaying, but immediately travels further than a radius of one (Figure 1.4(b)). It is thus similar to the cancer communication in that it has diffusion in all directions, but different in that each radius lingers and decays like the ant-like trail in the predator prey system, and in that the diffusion occurs immediately. It is also similar to both predator-prey communication paradigms as it is constantly sent, as opposed to relying on a specific event as in the cancer communication.

By investigating these four different approaches to communication for emergent behavior, I show how this type of communication can be utilized for a variety of goals. Additionally, I analyze each individual type of communication in its relevant chapter to determine the best parameters and how well it works for the designated task. This dissertation shows that communication based

on stigmergy can provide a simple yet effective way to increase collaboration among a variety of entities in a variety of systems.

CHAPTER 2

MAIA: MATHEMATICAL ANALYSIS OF CANCER THROUGH POPULATION DYNAMICS

2.1 Introduction

The necessity of including spatial consideration among entities, such as with the growth of trees, as well as how to adequately represent space has been a topic of debate in population dynamics modeling. Differential (ODE) and partial differential (PDE) equations have been utilized with both single and multiple populations, including for studying cancer cell growth. Differential equations generally either describe space globally or ignore it altogether, and can sometimes be analyzed and solved. Partial differential equations provide the most specific spatial information, but these models are often difficult or impossible to analyze without simulating the equations.

In this chapter we analyze the use of spatial requirements in studying the population dynamics of cancer cells and healthy tissue cells. The sets of healthy and tumor cells can be considered as two competing populations. They are cooperative within their own population set, in the sense that they are working toward the similar goal of having enough cells to fully populate the system. The cancer cells will be assumed to have some advantages over the normal cells, including a small probability that normal cells will mutate to cancer cells. Typically, models of competing populations assume globally bounded resources that limit the total number of units in the populations. We utilize the probabilities of creation and destruction of the cells to study local dynamics via approximation. Our two populations are not assumed to have identical spatial interpretations, and we also do not assume a specific dimensionality.

We develop four sets of ODEs that represent four different scenarios of spatial use by cancer and normal cells. Since ODEs are not always solvable, we use dynamical analysis to characterize the cell interactions for these scenarios. We evaluate how different views of space affect the growth of both types of cells, and in what situations cancer cells are most able to become the majority of cells in the system. Our results show that local considerations such as spatial requirements are crucial in

the study of evolutionary dynamics in cancer systems, and must be clearly postulated. Cancer cell ignorance of spatial requirements may be a crucial factor in cancer development. Although similar analysis has been done on cancer systems, we improve on those results by further analyzing the effect of spatial interactions on cancer and tissue growth.

Through the application of dynamical analysis to cancer cell interactions with normally functioning cells, we show that growth can be characterized by three fixed points and the final type of cell in the majority can be predicted. Although differential equation models of cancer exist, our model modifies the logistic equation to include varying types of spatial constraints based on both types of entities in the system, adding a dynamic between cell types that is usually ignored.

2.2 Related Work

The logistic equation is a classic tool in evolutionary dynamics for examining the growth of multiple different entities in a system [103]. This is not the first time an evolutionary dynamics approach has been applied to cancer growth [92, 103]. However, we expand on that work by modifying the logistic equation to include varying types of spatial requirements based on both types of entities in the system.

Hanahan and Weinberg put forth a number of ideas for the future of cancer research, such that research should be based on the set of “rules that govern the transformation of normal human cells into malignant cancers” [65]. These rules include proliferation, differentiation, and death. We propose an additional rule that is physical instead of biological: sensitivity to space requirements. Our dynamical system demonstrates the usefulness of this type of rule.

There are a number of ordinary differential equation (ODE) representations of cancer systems. Khain et al. utilizes diffusion coefficients in differential equations to represent nutrients in the system [78] and determine the dynamics of cancer cell growth in malignant brain cancer. These equations only examine cancer cells to determine the shape and speed of growth given different nutrient parameters. Sachs et al. surveys current ODE techniques for analyzing cancer cell growth [130]. We do not know of any ODE approaches that have cancer cells and healthy cells aware of each other, either spatially or otherwise.

Partial differential equations (PDEs) are also used to model cellular adhesion and density [57], nutrient location and dispersion [24], and the mutation to specific tumor suppressor genes [46].

Since PDEs of cancer growth often do not give as much extra information as one would like and become much harder to solve, we have chosen an ODE approach. Current PDE models focus on nutrients guiding tumor growth. However, as conjectured by Bru et al., tumor growth may be guided instead by space, and thus a spatial battle between the tumor cells and original cells occur [19]. For our model we implicitly take nutrition and growth hormones into account via probabilities, with the main focus on the interactions between tumor and normal cells as they vie for space.

In this chapter I will first discuss the series of equations in the ODE model of cancer growth. Then I will analyze these equations, discuss the results, and then conclude.

2.3 Model

Our model has two species: healthy tissue cells and cancer cells. Each cell type is modeled by the total number that exist over any given period of time. The number of cells increases by the addition of new cells to the system via replication; thus, new cells can only be added if at least one other cell of that type already exists. Cells can also be removed, and healthy cells can be converted to cancer cells. However, cancer cells cannot be mutated to healthy cells. All mutations are based on probability.

The resulting fluctuation of cancerous (PopC) and normal (PopN) cells over time is described by sets of differential equations on the number of PopC cells C , and the number of PopN cells N . These equations are based on probabilities: ρ_N , the normal cell proliferation probability; ρ_C , the cancerous cell proliferation probability; δ_N , the normal cell death probability; δ_C , the cancerous cell death probability; μ , the probability of mutation from normal cell to cancerous cell. The probabilities controlling our cell actions correspond to cellular functions and Hanahan and Weinberg's rules.

Although we do not explicitly model the gene interactions that cause mutation, proliferation, etc, these probabilities can capture the essence of those interactions. For instance, μ represents the probability that at any given point in time a cell will have acquired all necessary mutations to become a cancer cell. We do not include mutations returning a cancer cell to normal. Since we only model the genes implicitly we do not lose generality in the sense that we are not committed to any one theory of exactly how many genes are necessary. This further increases the possibilities of applying our method to multiple cancer types.

Defining growth of normal cells by the logistic equation would give us $\frac{dN}{dt} = \rho_N N * (1 - \frac{N}{Z})$, if Z represents the carrying capacity of our species N (i.e., maximum number of normal cells allowed in the system). The classic logistic equation does not use a rate of death, but instead modifies the proliferation rate based on how close the population is to the carrying capacity. Growth is initially exponential, but slows as the population reaches Z . We modify this equation by adding two explicit ways cells are removed from the population: death and mutation. We also modify the use of a carrying capacity to represent spatial requirements that may differ between our two populations.

We define four differential equations that represent four different spatial requirement situations (Table 2.1). The base pair of equation (Σ_1) has no spatial constraint, meaning that both PopN and PopC cells consider the world to allow endless development. At each period of time, the number of cells in the system increases based on the proliferation rate minus the death and mutation rates. There is no carrying capacity, or system maximum, to stop growth. Σ_1 thus adds death and mutation to the original logistic equation, but removes any consideration by cells of carrying capacity. This exponential growth model (Σ_1) acts as a base to compare to the other three equations.

To re-include the carrying capacity, we constrain the proliferation and death probabilities by the available space. A sensitive proliferation rate means that as the number of cells increases the proliferation ratio decreases; once the maximum count of cells exists the proliferation rate becomes zero. The spatial constraint on the death rate works in the reverse, with death rate increasing as the number of cells decreases. The spatial constraint is based on the ratio $0 \leq \Gamma \leq 1$ of the used space (Equation 2.1). If each PopN cell needs 1 unit of space, then a maximum of Z PopN cells can exist in the system with no PopC cells. We can allow PopC cells to require more or less space than PopN cells with a variable k , where each PopC cell uses k spaces as defined by PopN cells. If $k = 1$ both cell types require an identical amount of space, and if $k < 1$ ($k > 1$) then each PopC cell uses less (more) space than a PopC cell. Thus, we can represent the interactions between PopN and PopC cells in an environment of a defined size whether or not the two cell populations view space identically. The proliferation is multiplied by $(1 - \Gamma)$ so that when there is no available space the replication rate is 0. The death rate is multiplied by Γ so a lack of available space maximizes the death rate (δ_N or δ_C).

$$\Gamma = \frac{N + \frac{C}{k}}{Z} \quad (2.1)$$

Σ_1	$\begin{aligned} dN/dt &= (\rho_N - \delta_N - \mu)N \\ dC/dt &= (\rho_C - \delta_C)C + \mu N \end{aligned}$
Σ_2	$\begin{aligned} dN/dt &= (\rho_N(1 - \Gamma) - \delta_N - \mu)N \\ dC/dt &= (\rho_C(1 - \Gamma) - \delta_C)C + \mu N \end{aligned}$
Σ_3	$\begin{aligned} dN/dt &= (\rho_N(1 - \Gamma) - \delta_N\Gamma - \mu)N \\ dC/dt &= (\rho_C - \delta_C)C + \mu N \end{aligned}$
Σ_4	$\begin{aligned} dN/dt &= (\rho_N(1 - \Gamma) - \delta_N\Gamma - \mu)N \\ dC/dt &= (\rho_C(1 - \Gamma) - \delta_C\Gamma)C + \mu N \end{aligned}$

Table 2.1. The four equations representing different models of global spatial requirements.

The spatial considerations are based on the total number of cells instead of a cell's own type only, and both proliferation and removal can be affected. In Σ_2 both PopN and PopC cells change proliferation based on the amount of space available. In Σ_3 PopN cells also change death rate and PopC cells are blind to available space; in Σ_4 both types of cells are sensitive to change in both their proliferation and death rates. By analyzing these sets of equations we can determine the importance of space constraints on the final ratio of PopN and PopC cells.

2.4 Dynamical Analysis

We analyze these four equations to determine

1. **Minimum Death Value:** the minimum δ_C values that will lead to a normal cells majority versus a cancer cell majority, representing how difficult it is in each scenario to limit cancer cell growth.
2. **Prediction Equations:** when a specific number of PopN and PopC cells is likely or impossible to happen, how the variable k affects the system dynamics, and how the situations represented by $\Sigma_1 - \Sigma_4$ affect cell survivability.

In the first case we will simulate the equations in Mathematica, and in the second case we will solve a generalized version of the equations and utilize nonlinear dynamics analysis via fixed points.

2.4.1 Minimum Death Value

By simulating each set of differential equations $\Sigma_1 - \Sigma_4$, we can find relative parameter values that lead to a specific cell type as the majority over time (see Table 2.2). It is often useful to know the




Figure 2.1. Demonstration with Σ_3 and Σ_4 for $\rho_N = 0.01$, $\rho_C = 0.1$, $\mu = 0.00000001$, $\delta_N = 0.001$, $Z = 1000$. Solid lines represent the number of PopN cells, dashed lines represent the number of PopC cells. Changes in slope represent proliferation and death changes based on Γ . Each equation set ends in an equilibrium with PopN as the majority only when δ_C meets the minimum requirements for the equation set. For Σ_3 : (a) δ_C is too low, giving a PopC majority (b) δ_C is high enough to give a PopN majority. For Σ_4 : (c) δ_C is too low, giving a PopC majority (d) δ_C is high enough to give a PopN majority. Note that for Σ_4 PopC cells are constrained to the size of the system, whereas in Σ_3 they can increase infinitely.

type of cell that will eventually become the majority given no intervention, with majority defined as greater than 50% of the total system population. Understanding the role of spatial limitations on cellular growth will influence our interpretation of cellular situations and modeling decisions. Examples of the changes of cells over time can be seen in Figure 2.1. This figure shows Σ_3 and Σ_4 when δ_C is defined for a PopN majority, as well as for a PopC majority.

The four equation sets were simulated for $\rho_N < \rho_C$, as that is the hardest case to determine if the PopN cells will dominate. For $\rho_N > \rho_C$ the main determinant is whether or not $\delta_N + \mu > \delta_C$, as will be seen with Σ_1 . Each system of equations is simulated on 10 values for δ_N , three values of ρ_N , and all valid values of ρ_C for each ρ_N in increments of 0.1. For each set of simulations, the

	PopN Majority	PopC Majority
Σ_1	$\rho_N > \delta_N + \mu$ $\mu = 0$ $\delta_C \geq \rho_C$	$\rho_N \leq \delta_N + \mu$ $\delta_C < \rho_C$
Σ_2	$\delta_C \geq \frac{\rho_C}{\rho_N}(\delta_N + \mu) + \mu$	$\delta_C < \frac{\rho_C}{\rho_N}(\delta_N + \mu) + \mu$
Σ_3	$\delta_C \geq \rho_C + \mu$	$\delta_C < \rho_C + \mu$
Σ_4	$\delta_C \geq \frac{\rho_C}{\rho_N}\delta_N(1 + \frac{\mu}{0.05})$	$\delta_C < \frac{\rho_C}{\rho_N}\delta_N(1 + \frac{\mu}{0.05})$

Table 2.2. Approximation of minimal removal values for all equations based on results from simulation of the ODEs.

minimum δ_C value was determined to three decimal places for each δ_N such that PopN cells end as the majority.

Given the simulation data we can fit a curve to produce a set of equations that approximate the minimum death value needed to provide either a healthy or cancer cell majority. The results shown in Table. 2.2 were determined by approximating these simulation curves in Mathematica. All equations fit the results with R^2 values of at least 0.99, with a value of 1 denoting a perfect prediction.

2.4.2 Prediction Equations

Now that we have a basic understanding for how the equations relate, we analyze their fixed points. A fixed point is a value of N and C that affects the equation dynamics. This analysis will show when a specific number of PopN and PopC cells is likely or impossible to happen, how the variable k affects the system dynamics, and how the situations represented by $\Sigma_1 - \Sigma_4$ affect cell survivability.

$$\begin{aligned}
dN/dt &= N(A_1N + B_1C + D_1) \\
dC/dt &= M(A_2C + B_2C + D_2) + \mu N
\end{aligned} \tag{2.2}$$

After generalizing the equations to the nonlinear forms in Equation 2.2, we locate the fixed points. This general form allows us to analyze only one set of equations and then apply the results to all of our original four sets of equations based on the definitions of the parameters in the equations in Equation 2.2 in terms of each original equation set. There are three fixed points for these equations that were found by setting the generalized equations to 0, as seen in Table 2.3.

	fx_1	fx_2	fx_3
N	0	0	$-\frac{1}{A_1}(B_1C + D_1)$
C	0	$-\frac{D_2}{B_2}$	$\frac{\mu D_1}{A_1 D_2 - A_2 D_1 - \mu B_1}$
λ_1	D_1	$D_1 - \frac{B_1 D_2}{B_2}$	see text
λ_2	D_2	$-D_2$	see text

Table 2.3. Fixed Points and their eigenvalues for the general form of the equations.

The real parts of the eigenvalues (λ_1, λ_2) related to each fixed point are analyzed to determine when the fixed point becomes an attractor ($\lambda_1, \lambda_2 < 0$), repeller ($\lambda_1, \lambda_2 > 0$), or neutral ($\lambda_1, \lambda_2 = 0$). A fixed point can also become a saddle point when one eigenvalue is negative and the other is positive. When a fixed point is acting as an attractor, if the system enters a nearby state it will move toward the fixed point. When it is instead a repeller, the opposite effect will occur. We classify a fixed point as neutral in situations where it will neither attract nor repel the values from neighboring states; it essentially is as if it does not exist. A saddle point will both attract and repel.

All fixed points can act as attractors, repellers, neutral, or saddle points, but for any set of equation parameters they will each be only one of the four. The fixed points we find for our general equations can be redefined for each of our initial equations in terms of our initial parameters $\rho_C, \rho_N, \delta_C, \delta_N, \mu$ as seen in Table 2.4 and 2.5.

The phase planes of these fixed points (Figure 2.2) demonstrate growth patterns for each Σ . A point on the phase plane shows the total number of PopC cells (x-axis) and the total number of PopN cells (y-axis). The larger dots represent the fixed points from Table 2.3 that are not neutral or negative for those parameter values. To determine the changes that will occur for the two populations, choose a starting point (number of each cell type) and follow the arrows. The total will either end at an axis, going off the chart, or at a fixed point. Different starting points can potentially lead to different ending points for the same parameter set. For instance, only the fixed point fx_1 is valid for Σ_1 . For Σ_2 any of the three fixed points may appear for any given parameter set, and Figure 2.2 shows only fx_1 and fx_2 as active. Since fx_2 is not valid for Σ_3 , Figure 2.2(c) only has two fixed points corresponding to fx_1 and fx_3 . However, all three fixed points are possible for Σ_4 , and are all shown in the phase plane.

Analysis by each Σ

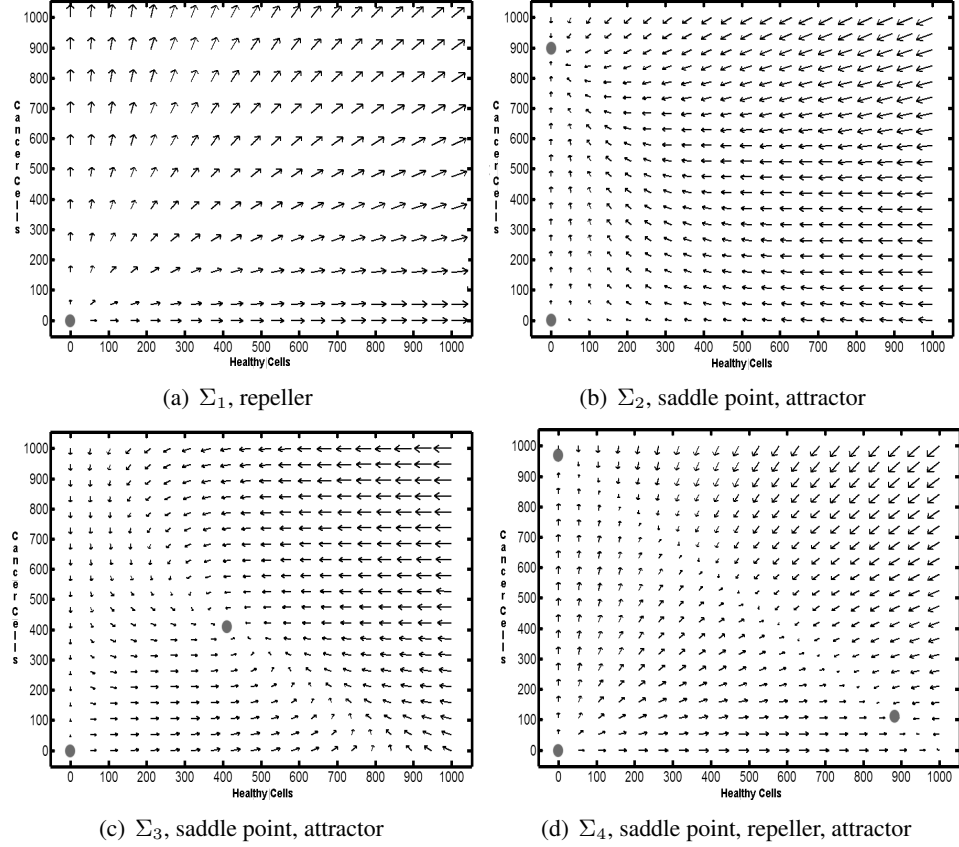


Figure 2.2. Phase planes for the four equations and different parameter values, with red circles representing fixed points. (a) $\rho_N = 0.1, \rho_C = 0.1, \delta_N = 0.01, \delta_C = 0.01, \mu = 0.001$ (b) $\rho_N = 0.1, \rho_C = 0.1, \delta_N = 0.1, \delta_C = 0.01, \mu = 0.01$ (c) $\rho_N = 0.1, \rho_C = 0.1, \delta_N = 0.01, \delta_C = 0.11, \mu = 0.01$ (d) $\rho_N = 0.3, \rho_C = 0.3, \delta_N = 0.001, \delta_C = 0.01, \mu = 0.001$

Σ_1 is the most simplistic, as only fx_1 applies. It thus only has two options: exponential growth, or complete removal of all cells of all types. All cells die if the PopN removal and conversion rates sum to more than the replication rate, and the PopC removal rate is higher than the PopC replication rate. The opposite relationships will cause all cells to grow to infinity. This can be seen in the first row of Table 2.4.

When only the replication rates are affected by space (Σ_2) it is also true that all cells die if the PopN removal and conversion rates sum to more than the replication rate, and the PopC removal rate is higher than the PopC replication rate (second row, Table 2.4). However, the cells will not necessarily grow to infinity due to the opposing relations. Instead, one of the other fixed points may become relevant. PopN cells may die out with PopC retaining the number of available spaces based on the ratio of removal to replication rate (fx_2). This fixed point is only an attractor when

	Attractor	Repeller	Neutral
Σ_1	$\rho_N < \delta_N + \mu$ $\rho_C < \delta_C$	$\rho_N > \delta_N + \mu$ $\rho_C > \delta_C$	$\rho_N = \delta_N + \mu$ $\rho_C = \delta_C$
Σ_2	$\rho_N < \delta_N + \mu$ $\rho_C < \delta_C$	$\rho_N > \delta_N + \mu$ $\rho_C > \delta_C$	$\rho_N = \delta_N + \mu$ $\rho_C = \delta_C$
Σ_3	$\rho_N < \mu$ $\rho_C < \delta_C$	$\rho_N > \mu$ $\rho_C > \delta_C$	$\rho_N = \mu$ $\rho_C = \delta_C$
Σ_4	NONE	$\rho_N > \mu$ $\rho_C > 0$	$\rho_N = \mu$ $\rho_C = 0$

Table 2.4. Fixed Point f_{x_1} ($C = 0, M = 0$) for each equation type, written in our original variables

	Σ_2	Σ_4
$\frac{-D_2}{B_2}$	$\frac{\delta_C}{\rho_C} kZ - 1$	$\frac{\rho_C kZ}{\rho_C + \delta_C}$
Attractor	$\frac{\delta_C}{\rho_C} < \frac{\mu + \delta_N}{\rho_N}$ $\delta_C < \rho_C$	$\frac{\rho_C(\rho_N + \delta_N)}{\rho_C + \delta_C} > \rho_N - \mu$ $-\rho_C < 0$
Repeller	$\frac{\delta_C}{\rho_C} > \frac{\mu + \delta_N}{\rho_N}$ $\delta_C > \rho_C$	NONE
Neutral	$\rho_N = \mu + \delta_N$ $\delta_C = \rho_C$	$\rho_N - \mu = 0$ $\rho_C = 0, \delta_C > 0$

Table 2.5. Fixed Point f_{x_2} for each equation type, written in our original variables. This fixed point is only valid for these two cases.

the removal to replication ratio for PopC cells is less than the removal and conversion to replication ratio for PopN cells (first column, Table 2.5). It is harder to quantify when f_{x_3} will occur, but a few examples are in Table 2.6.

The Σ_3 model is only affected by f_{x_1} and f_{x_3} . All cells will die if the PopN replication rate is less than the conversion rate, and the PopC replication rate is less than the PopC removal rate (third row, Table 2.4). Since only the PopN cells are affected by space constraints, it is possible for the PopC cells to grow to infinity while the PopN cells die out. This is one way to reason about why f_{x_2} is not relevant for Σ_3 .

The most complicated of the models, Σ_4 , is affected by all three fixed points. The fixed point f_{x_1} cannot attract in this case, however it does repel both cell types when the PopN replication rate is more than the conversion rate, and the PopC replication rate is greater than zero (fourth row, Table 2.4). Unlike the other equation sets, neither removal rate is relevant. The second fixed point f_{x_2} cannot repel (second column, Table 2.5).

Analysis By Each Fixed Point

We also analyze the fixed points individually. The trivial fixed point fx_1 (Table 2.4) represents the removal of all cells in the system. For PopN cells modeled by Σ_1 or Σ_2 , if they do not replicate faster than they convert and die, fx_1 is an attractor (all cells die); if they replicate faster, then fx_1 is a repeller (cannot remove all cells from the system); and if they replicate at the same speed, fx_1 is neutral. The same is true for PopC cells except that they compare replication with removal rate only. If only PopN cells have space constraints (Σ_3), the PopN removal rate is no longer relevant. This is also true for Σ_4 , except that the fixed point cannot be an attractor.

The fixed point fx_2 is only valid for Σ_2 and Σ_4 (Table 2.5). For Σ_2 the eigenvalues relate the ratio of PopC cell removal to PopC cell replication, and the ratio of PopN cell removal (which includes conversion) to PopN cell replication. Thus, it is an attractor when PopC cells have the higher ratio and they replicate more often than they die, and it is a repeller when PopN cells have the higher ratio and PopC cells die more often than they replicate. It is neutral when PopC cells die and replicate at the same rate, and PopN cells replicate at the rate that they die and convert.

The eigenvalues for fixed point fx_3 are more complicated, and thus will not appear explicitly. However, we show examples of this fixed point with different sets of variables in Table 2.6. For many valid variable values fx_3 is defined with one or both cell type quantities negative, which is not a valid state. Thus, this fixed point is not often relevant, but it is generally an attractor when it is valid.

The variable k (factor of space difference between PopC and PopN) does not appear in any of the eigenvalues for the three fixed points. This shows that the aggressiveness of the PopC cells as described by how much less (or more) space they take than PopN cells does not affect the end result of the system. This result was also found in the Minimal Death Value analysis. In this type of model at least it is not important to consider that cancer cells can exist in a denser grouping than healthy cells. This result should be further examined in other systems.

It is possible for each fixed point to become a saddle point. A saddle point at the origin indicates a situation where the cells do not die out. This occurs when one cell type dies faster than it replicates, and the other cell type replicates faster than it dies. For PopN cells, the total removal rate includes the conversion rate. Thus, if one cell type dies out the other one will thrive.

	ρ_N	ρ_C	δ_N	δ_C	μ	$PopN$	$PopC$	λ_1	λ_2
Σ_2	0.3	0.3	0.001	0.01	0.001	996.793	-3.43722	-0.299526	0.289491
Σ_3	0.3	0.3	0.001	0.01	0.001	882.963	110.37	-0.298	-0.008
Σ_4	0.3	0.3	0.001	0.01	0.001	882.244	111.111	-0.300027	$-7.913e^{-3}$
Σ_2	0.3	0.3	0.001	0.4	0.001	983.52	9.8352	-0.294518	-0.101522
Σ_3	0.3	0.3	0.001	0.4	0.001	990.844	2.48956	-0.398	-0.298
Σ_4	0.3	0.3	0.001	0.4	0.001	990.849	2.50627	-0.399329	-0.29602
Σ_2	0.2	0.9	0.001	0.01	0.001	991.163	-1.11367	0.889817	-0.199041
Σ_3	0.2	0.9	0.001	0.01	0.001	495	495	-0.545137	-0.000363
Σ_4	0.2	0.9	0.001	0.01	0.001	481.098	508.951	-0.560456	$-3.356e^{-4}$
Σ_2	0.001	0.9	0.01	0.001	0.001	0	0	0.899	0
Σ_3	0.001	0.9	0.01	0.001	0.001	1.01031	-10001	9.89809	0.01
Σ_4	0.001	0.9	0.01	0.001	0.001	0	0	0.9	0

Table 2.6. Examples of fixed point values for different parameter values for the fixed point fx_3 when $k = 1$ and $N = 1000$. These values are obtained via simulation of the equations.

2.5 Conclusions

By analyzing these four sets of equations we see that spatial limitations for both cell types will impact the final result of cell growth and sustainability. No spatial limitation results in exponential growth (Σ_1), and spatial limitations on both proliferation and death rates limits growth the most (Σ_4). If only PopN cells consider space then a higher PopC death rate is necessary for the PopN cells to become the majority (Σ_3); the PopC cells will also be less likely to die out in general. However, for all types of space requirements (Σ_2 to Σ_4) we are able to predict the type of cell that will end in majority, and if either or both cell types will die out.

Although initial intuition points to exact spatial representations, an approximation of the space can be both advantageous and equally effective. Utilizing probability can result in good predictions of the ratio between the different cell types, without the cost of analyzing partial differential equations. It can also allow the researcher to focus on higher level issues, without delving in to specifics that are not always known to be entirely true but are part of a hypothesis themselves. The logistic equation is similar to our equations but does not address issues of global space in the same way. We expand on this classic evolutionary dynamic approach by requiring all cell types to consider all other types when determining available space, as well as affecting both the proliferation and removal rates by spatial constraints.

We are thus able to predict the interactions between malicious and cooperative cells on a global scale. The fixed points describe when specific types of behaviors will occur: the eradication of all cells in the system, the growth of all cells to a specific total, and the destruction of one cell type in favor of another. The simulation results give further insight into when we can expect a cooperative majority in the system, which is surprisingly often.

These results can directly influence future models, including our choice to use a three-dimensional agent-based model to continue evaluating cancer dynamics in the next chapter. Utilizing a three-dimensional model is one way to directly model spatial limitations, as opposed to using the global view in the ODE model. Our results imply that healthy cells can often become the majority in a tissue without outside intervention. We acknowledge that different intra-cellular processes cause the cellular actions, but ignore that detail in favor of learning about the end picture. Our equations thus appear able to predict the ratio of healthy to cancer cells with minimal knowledge of the initial situation, and may be used in the future to classify cancers based on their growth rates more readily than is allowed by current models. They also imply that explicitly being aware of the spatial needs of cells, including the nutrient availability, is crucial to an accurate model of cancer. This implication is based on the fact that different versions of the equations, with varying use of explicit space, result in different numbers of healthy and cancer cells.

CHAPTER 3

AGENT-BASED CANCER AND COMMUNICATION MODEL

3.1 Introduction

We introduce an original computational system to identify logical principles underlying cancer development and suggest innovative anti-cancer solutions. The model abstracts the complex environment of cancer development and progression, where numerous chemical, biological, and physical factors act together to affect intra-cellular events and extra-cellular signaling. While simplification is essential to unmask the fundamental principles of cancer occurrence, the artificial intelligence component of our system affords a high level of adaptation for numerous intra- and extra-cellular details, unlike previous cancer models that were restricted to probabilities of several intra-cellular events [84, 101, 142, 129, 54, 139, 53, 7]. Our model provides a framework to assess several important questions in Oncology: What kind of information flow inside and between cells may be associated with tumor development and progression; What kind of inter-cellular communication keeps tumor cells dormant; Do current therapies bias some of the natural flow to explain their temporary benefit; And what are the principles of successful inter-cellular communication rules that would enable selective tumor cells' apoptosis (programmed cell death). The latter subject is the focus of this chapter.

We relate two seemingly opposing biological facts about cancer and apoptosis: The classic hallmark of cancer, that cancer arises when inappropriate apoptotic response occurs and prevents natural eradication of mutated cells [65], and the induction of caspase-dependent tumor cell apoptosis as a universal mechanism for tumor cell death by irradiation and the majority of chemotherapy agents [123, 52, 41]. This leads to our assumption that even highly mutated cells maintain residual apoptotic abilities.

Our model seeks to provide a single unifying mechanism incorporating these assumptions [9, 69, 121, 10]. It is a dynamic tissue simulation model composed of cells cycling in a 3-dimensional

society, where normal and mutated cells are defined phenotypically by their apoptotic response, proliferation pace, and compliance with spatial regulation rules. Signaling toward apoptosis is modeled in terms of the flow of information that activates the intrinsic and extrinsic paths correlating with activities of intrinsic regulating factors, and extrinsic modifiers, respectively [69, 121]. Findings suggest that our modeling may constitute a complementary approach to biological research; directions proposed by the model could lead to an enhanced understanding of these processes and of potential interventions.

Tumor growth is monitored by many processes, including the diffusion of nutrients. A tumor cluster will slow its growth as nutrients become scarce. As nutrients are depleted by the tumor, the cells that do not receive enough nutrients to function normally will end up in one of two different states: necrotic or hypoxic. A hypoxic cell receives enough nutrient to continue surviving, but not enough to proliferate. A cell may remain hypoxic indefinitely. A necrotic cell is a cell that has received even fewer nutrients, such that it does not receive enough nutrients to survive. Once a cell becomes necrotic, there is no way to revive it; it will eventually die.

Tumor cells, however, are often able to convince nearby vasculatures to grow such that the tumor can receive enough nutrient to continue growing. This process is called angiogenesis, and is accomplished by hypoxic cells emitting chemicals into the environment. Once these chemicals meet a vasculature, they can encourage the growth of endothelial cells for the creation of new vasculatures. These sprouts will only provide nutrient to the tumor once they have created a loop so that blood can flow. An overview of angiogenesis is provided by [89].

Our goal is to build a three-dimensional model of tumor growth that includes normally functioning tissue cells and tumor-induced angiogenesis, to examine the case of cancer cell death. Although cancer cells have lost their initial ability to undergo apoptosis, there are still available secondary mechanisms to activate self-death. Natural mechanisms for selective cancer cell death are supported by the fact that otherwise cancer incidence as calculated by mutation rates would be significantly higher [101]. Analogous mechanisms have been reported and include tumor removal via immune surveillance [74].

We hypothesize that one mechanism for cancer removal requires communication with other cells in the system, both by signal emission on viability status (alive or dead) and the compliance with external apoptotic commands. An example of communication playing a role in cancer death is the

role of high mobility group box 1 (HMBOX1) protein in reporting cancer cell death to the immune system as an essential part of tumor death by chemotherapy agents [9]. Findings also show that AP2L/TRAIL (tumor-necrosis-factor-related apoptosis induced ligand) binds to “death receptors” DR4 and DR5 to induce selective tumor cell apoptosis via the external apoptosis pathway [69, 121, 10]. This also supports the idea that cancer cells can still die, and may die from communication.

This chapter presents a new agent-based cancer model that includes both cancer and normal tissue cells interacting in three-dimensions. It has been suggested that three dimensional models have the most to offer for bridging the gap between in vivo studies in two-dimensional cultures and the larger animal system [154]. Agent-based models allow each cell to be modeled separately, following some series of rules on its behavior. Blood vessels and angiogenesis are also included to limit growth based on nutrient availability. We analyze the role of inter-cellular communication in the removal of cancer cells, based on the fact that cells may die due to neighbors undergoing apoptosis. Through the model we hypothesize that this type of communication may explain the natural removal of cancer cells within the human body, and examine in what cases this type of communication mechanism would succeed in cancer removal. In the next section we will discuss additional related work, then describe the cancer model, followed by a description of the communication paradigms, results, and then conclusions.

3.2 Related Work

3.2.1 Spatial Cancer Models

There are many valid techniques for studying cancer, as discussed in Chapters 1 and 2. As was shown in Chapter 2, the choice of spatial vs. non-spatial model (such as an ODE) model can have a huge impact on the results. We choose to utilize a spatial representation of our cells because it is more similar to the original biological system and thus allows for a greater number of detailed aspects to be included. To explicitly represent space the main options are partial-differential equations (PDE), a cellular automata model, or an agent-based model. We choose to use agent-based models as they provide more freedom in development, and are a great fit for modeling cancer as it creates a non-linear system where individual cells not only determine their own next step but can easily interact with other cells. They also have the benefit of being able to be simulated in a more accurate way than how one would simulate a PDE. We investigate the use of agent based models

to represent both normal and cancer cells, whereas models generally only represent cancer cells explicitly.

We develop a three-dimensional agent-based model, although two-dimensional models have also been studied. CancerSim is a three-dimensional agent-based model that begins with a cancer-free system and then focuses on the growth of the cancerous cells using the hallmarks of cancer [65] as its basis [2]. It therefore takes into account basic cellular properties such as genetic instability, telomere length, and random apoptosis. Another three-dimensional agent-based tumor model utilizes basic gene-protein interactions and multi-cellular patterns specific to brain cancer [156]. This model represents internal cellular processes via differential equations, and the location of cells spatially.

Cellular automata models have also been developed to specifically model cancer cell growth. Quaranta et al. develop a two dimensional model to study tumor spheroid growth based on reaction-diffusion-taxis equations, with a plan to move the model to three dimensions [122]. Gerlee et al. build on this model by adding angiogenesis and controlling each cell with a neural network [58]. Cellular automata have been used to model invasive cells as a wave of invasion, determining how a carrying capacity can help correctly model this type of growth [139].

None of the above models consider how the cancer cells affect the rest of the system, however. One way to involve other cells is by investigating the role of the immune systems. For instance, an agent-based approach is used to investigate the role of the immune system in reacting to and fighting tumor growth [43]. We design our model to also include both cancer and healthy cells, although we include healthy tissue cells instead of the immune cells. Including other types of cells leads to a more realistic model, as surrounding cells have already been shown to play a role in affecting cancer growth.

3.2.2 Angiogenesis Models

Models of angiogenesis fall in to three different overall categories: models of vasculature growth without tumor or other tissue cells, models of vasculature growth with only tumor cells, and models of vasculature growth with both tumor and healthy tissue cells. Generally, models of angiogenesis are attempting to model the growth of vasculatures to support tumors and the size of tumors with and without angiogenesis. In some models angiogenesis is just one of the mechanisms modeled in

order to answer a different question, such as the shape of growth or time of growth of the tumor. This is the type of model we create. Both [113] and [26] provide a good overview of the process of angiogenesis, as well as the state of the art of angiogenesis modeling.

Various levels of mathematical modeling exist to include how capillaries are formed, the flow of blood, vessel adaptation, and the extent of which chemical diffusion reaches the tissue [94]. McDougall et al. also analyze how well drugs could be sent to the tumor via the formed vessel network, based on the quality of the created vessels. They found that the simulation was sensitive to changes in parameters for haptotactic response of the blood vessel cells, blood viscosity, and blood pressure, and that the tracer-drug sent through the system was able to travel easiest through the well formed vessels.

A discrete mathematical model based on PDEs has also been used to examine the spatio-temporal evolution of capillary networks in two and three dimensions [25]. They initialize with a vessel at one end of their model and a tumor at the other end, a gradient of TAF between the tumor and the vessel, and five sprouts started at the vessel. Capillaries then grow by following the gradient of TAF. They find that the model provides capillaries with realistic structure and morphology.

A two-dimensional multi-scale model of angiogenesis where endothelial cell sprouting occurs due to VEGF diffused by healthy cells in an initial development phase, and by tumors that are later implanted into the model, has been developed by [111]. Sprouts grow based on VEGF gradient that occurs when cells do not have enough oxygen. Sprouts die if they do not create a loop within a given period of time, and are only active once it is a viable segment. They find that a larger number of tumor cells causes a higher vascular density, and that network remodeling requires a balance between angiogenesis and vessel pruning. A two-dimensional PDE model of tumor growth with angiogenesis has also been developed [137]. They find that asymmetrical tumor growth leads to a greater degree of branching at the surface of the tumor when compared to symmetrical tumor growth.

Multiple models are based on the cellular Potts model, a freely available model that gives a basis for cellular modeling that includes PDE solvers to determine diffusion for nutrients or other user defined chemicals such as VEGF. [12] create a cell-based model of angiogenesis based on the cellular Potts model, and find that the sprout morphology is not affected by how far from the tip the proliferating region is, although the speed of growth is affected. [138] presents a three-dimensional

simulation of tumor growth and angiogenesis based on the cellular Potts model. They analyze tumor growth with and without angiogenesis. The initial vasculature is represented as a grid. Endothelial sprouts grow in biased random movement toward VEGF secreted by hypoxic tumor cells, with the only modeled nutrient being oxygen. They find that the system goes through six phases of growth, and that their results are similar to other models.

Cellular automata models have been used to investigate branching properties of vessels [91][149][114]. Some models focus on branching rules and how they change the vasculature, with a brief description of how it may be applied to tumor angiogenesis [91]. Others focus specifically on tumor angiogenesis, with either a visual comparison to in vivo endothelial sprouting [149] or a comparison with data on in vivo vascular growth [114]. Both find that their sprouting is similar to what would be expected from tumor angiogenesis.

Agent-based models are currently underdeveloped in cancer biology, both for general growth models as well as angiogenesis models. Agent-based models enables the combination of many different levels and types of detail and interaction, which may be difficult in more constrained systems such as ODEs, PDEs, or cellular automata. However, there are not many tools available for non-programmers to develop these types of models. Also, when they are developed it can be difficult to design them such that they can run in a reasonable amount of time given the resources a researcher is likely to have available. However, it is possible to design algorithms to save time and still provide the correct output for this type of complex system. The model described in this chapter is an agent-based model designed to incorporate many aspects of cancer growth, but with algorithmic choices that allow a complex model to still be able to be run in only a few hours.

3.3 Cancer Model

In our model, each cell exists on a three-dimensional grid to approximate a tissue. Although two-dimensions can also be modeled, three dimensions provides a more realistic representation. Cells follow individual “life protocols” defined and implemented through probabilities represented within their computational genes. We consider normal tissue cells to be of size $10\mu m$, and each time step to represent approximately one day. Vasculatures in the system give cells nutrients, which they require to survive. Without tumor growth, the system can self-maintain essentially forever.

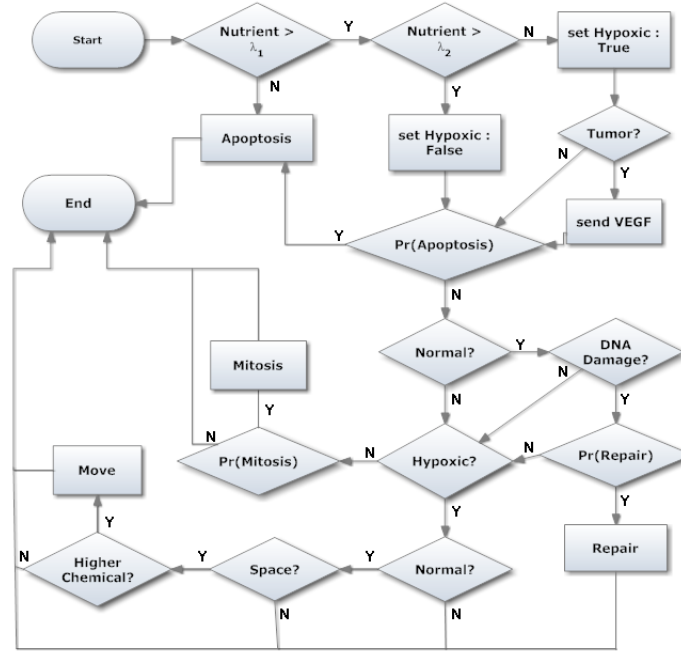


Figure 3.1. Flow chart for both normal and cancer cells, for every time step. The differences for normal and cancer cells are what probabilities are used and how the genes affect decisions (not shown).

3.3.1 Genes and the Life Protocols

The basic life protocols reflect proliferation (including rate parameters, generation potential, and space restrictions [6, 128, 11]), proliferation-suppression mechanisms, self-testing at a check point prior to the replication decision, repairing damage, and apoptosis (self-death). The latter is activated as either a random process, secondary to a cell's decision to die due to aging or uncorrected defects, or as a reaction to extra-cellular signaling. The distance regulation protocol maintains shape cohesiveness and allows undisturbed communication flow among cells [7].

These protocols are chosen to approximate the healthy functioning of phenotypic properties of cycling cells. The physical property of space is important in our model since normal cells proliferate only when given space around them, whereas tumor cells may violate this restriction. This can be a basis for the creation of solid tumors of a particular shape, as our cells grow in an expected spheroid pattern with current spatial parameters even though we do not explicitly model this behavior [156, 59]. They also relate to the hallmarks of cancer, which are generally used as a basis for cancer models [65]. A flow chart of the decisions by a cell at every time step can be seen in Figure 3.1.

Tumor cells develop in the system when all of a cell's life protocols are damaged. A tumor cell can only produce tumor cells and cannot back mutate into a normal cell. Since all life protocols are broken, a tumor's ability to follow each of those protocols is also broken: it is not able to repair the mutations, it cannot undergo apoptosis that would occur due to mutations, and it proliferates more frequently than a healthy cell. Tumors are able to exist eight times denser than normal cells, thus for the amount of space saved for a normal cells (including space buffer) it is possible for eight tumor cells to exist within that same area.

3.3.2 Nutrients

All cells require nutrients to survive, which they acquire from blood vessels. We specifically model oxygen as it is generally the most limiting nutrient sent to cells. Each blood vessel diffuses nutrients to nearby cells; there is a limited amount of nutrient that can be supplied by a single vessel, and a limited distance that it can travel from the vessel. Additionally, as cells use nutrient it decreases the amount of nutrient available for cells further from the vessel. Thus, vessels must occur close enough to sustain all cells, but optimally should be as far apart as possible to accomplish that goal.

We represent blood vessels at equidistant locations throughout the tissue (every $200\mu m$ since oxygen can diffuse up to $100\mu m$), forming a grid. We further view the grid as creating a series of compartments, with each compartment being bounded on three sides by vessels, one in each of the three directions. These vessels can supply nutrients to all healthy cells within the tissue, assuming there are no tumor cells present. Nutrients are distributed to cells from each vessel, first to the cells closest to the vessel, then to the next layer of cells past that, etc, until the maximum number of cells that vessel can support has been reached. We assume blood flows from the zero location to the maximum location for each vessel. Each vessel also has a maximum number of cells that can be supported overall, in addition to the maximum amount within each region.

To decrease the complexity in the model, instead of explicitly modeling the diffusion process we model the effect of diffusion. Usually diffusion is modeled using PDEs that must then be solved to determine the amount of gradient available at every location within the grid. Even discretizing this using standard grid diffusion techniques becomes computationally expensive due to the fact that the amount of nutrient available to a cell not only relies on the amount of nutrient in its location, but also

Algorithm 3.1 Diffusion of nutrients

```
1: compartmentNutrientMax
2: vesselNutrientMax = compartmentNutrientMax * numberOfCompartments
3: for each compartment in vessel do
4:   availableNutrient = compartmentNutrientMax
5:   for each diffusion radius do
6:     for each location over length of compartment at this radius do
7:       if availableNutrient < 0 then
8:         availableNutrient = 0
9:       end if
10:      send availableNutrient to each of N cells at location
11:      availableNutrient = availableNutrient - N
12:    end for
13:  end for
14: end for
```

on how many cells have taken nutrient from the environment before the nutrient reaches that cell. Thus, the diffusion also relies on the number of cells between a cell and its closest vessels. To take this issue into account, we instead do a stepwise diffusion from each vessel as seen in Algorithm 3.1.




Figure 3.2. Cells cycle through the normal and hypoxic states, depending on the amount of nutrients they receive. Once they are not receiving nutrients they become necrotic and die.

Cells that receive enough nutrients are considered normal and function as described previously. It is expected that cells will receive nutrients from three vessels: one in the X direction, one in the Y direction, and one in the Z direction. Thus, the minimum amount of nutrient for a cell to remain normal (λ_1) is three times what each vessels should give them. Cells that receive nutrients but not enough to remain normal (λ_2) will become hypoxic, and thus unable to proliferate until the nutrients increase again. Cells that receive zero nutrients will become necrotic and die immediately. These rules apply to both healthy and tumor cells, and are summarized in Figure 3.2.

3.3.3 Angiogenesis

Angiogenesis is the creation of new blood vessels. New blood vessels are formed during normal development, but hypoxic tumor cells can also induce vessel growth through angiogenesis. Without angiogenesis, a tumor's development will be limited, so angiogenesis is necessary to sustain continued growth. Additionally, angiogenesis is necessary to properly model the growth of the tumor, as otherwise growth will be faster than is realistic.

In the model, hypoxic tumor cells can diffuse VEGF to induce angiogenesis (Figure 3.1). VEGF is diffused from all hypoxic cells, and once the chemical reaches a vessel it will induce the creation of a sprout if there is not already one within a specific distance. Each sprout can grow by 5 endothelial cells every time step, and we assume that endothelial cells are half the size of healthy cells. Sprouts grow by biased random movement based on the gradient of VEGF within the system. They are also limited in that they can only grow in locations where other cells are not already present, and they are not allowed to form loops with themselves.

A sprout does not provide nutrients to surrounding cells until it forms a loop with another sprout or vessel. Once the sprout becomes viable, it will provide nutrients to nearby cells but not to the same extent that a vessel could provide. This limited nutrient, however, will allow the tumor to grow further, and create more sprouts as it grows as well.

3.3.4 Model Validation

First we analyze the growth of tumor and healthy cells together, with and without angiogenesis (Figure 3.3). As expected, cancer cells are able to grow much larger with angiogenesis, and normal cells are essentially removed from the system. We next validate the model by running the simulation with only tumor cells, as our comparison model does not contain normal cells. We start with a single tumor cell, which eventually grows into a tumor that can no longer be sustained by the vasculature. The size of this tumor is as would be expected. Additionally, the tumor grows in a spheroid pattern with small branchings, which is also expected.

Once the vasculature can't sustain tumor growth, hypoxic cells diffuse VEGF and angiogenesis begins. We examine the rates of tumor cells, hypoxic cells, and neovascular cells in this model (Figure 3.4) over time. The rates of tumor growth with and without angiogenesis relate as we would expect, with tumor growth increasing much faster with angiogenesis. Additionally, we see




Figure 3.3. Change over time in number of cancer (red) and healthy tissue (blue) cells. (a) When angiogenesis is not activated, cancer cells grow quickly, but are unable to grow large enough to crowd out healthy cells. (b) When angiogenesis is activated, the cancer cells are able to grow much larger and take nutrients and space away from healthy cells such that they basically disappear.

an increase in hypoxic cells in both cases. Eventually we start to see a more exponential growth in the case of angiogenesis, showing that the tumor is able to grow but still cannot support all cells. These graphs are as we expect to see, and are very similar to the graphs found in [138].

What we see in the growth of cells within the model when both normal tissue cells and cancer cells exist, is that the cancer is able to grow to a large percentage of the system. This result occurs whether or not angiogenesis is used within the model. For the rest of the chapter we will only consider a system in which both normal and cancer cells exist, and where cancer cells are able to influence angiogenesis.

3.4 Communication Model

Since cancer cells are able to grow such that they may overrun the system, we investigate methods of intercellular signaling that could result in the removal of cancer cells from the system. It has been seen that natural mechanisms for removing cancer cells should exist, otherwise cancer incidence would be much higher [101].

Our model includes intercellular signaling that occurs via diffusion of messages away from the sender. We develop a double messaging system with two forms of signals that can be sent by cells, both of which over time induce apoptosis within the receiving cells. It is understood that cancer cells still maintain the ability to undergo secondary apoptosis. For instance, AP2L/TRAIL (tumor-




Figure 3.4. Change over time with (blue) and without (red) angiogenesis: (a) Number of tumor cells, (b) Number of hypoxic (low nutrient) cells, (c) Number of new vessel cells, and (d) Number of non-hypoxic tumor cells. More tumor cells are able to grow with angiogenesis, and eventually there is also more rapid creation of hypoxic cells with angiogenesis. These graphs are very similar to those in [138].

necrosis-factor-related apoptosis induced ligand) binds to death receptors DR4 and DR5 to induce selective tumor apoptosis via an external apoptosis pathway [69, 121, 10].

The first message is called I'M DYING. This signal represents the scenario of cells dying due to their neighbor dying, and can affect both normal tissue and cancer cells. Each time a cell dies due to a specific reason, it will first send the I'M DYING message into the environment, and then it will undergo apoptosis. A cancer cell dying is more likely to be surrounded by cancer cells than by normal tissue cells. This process may be similar to HMBOX1, high mobility group box 1 protein, that is shown to be part of reporting tumor death by chemotherapy agents [9].

For the I'M DYING message to work there must be some mechanism to cause it to occur. We thus hypothesize a second signal that we call PLEASE DIE. This message will only be sent to neighbors when a cell is pushed by another cell. A push is defined as either a cell moving to a location where




Figure 3.5. The PLEASE DIE signal is first sent due to a push occurring, as seen in the middle image of (a). Once a cell is ready to undergo apoptosis from these signals, it will send the I'M DYING signal as seen in the middle image of (b), and then will undergo apoptosis as seen in the last image of (b).

another cell already exists or proliferating into a location already occupied by another cell, thus causing that other cell to move to an adjacent location.

These two message types thus work together to attempt the removal of cancer cells. The PLEASE DIE message reacts directly to an anomaly in the system that is in this case based on spatial availability. The I'M DYING message can then propagate the message further, as once a cell is dying from the messages it is likely that it was a cancer cell. The interaction of these two messages can be seen in Figure 3.5.

Both of the signals have a strength and a distance. The distance defines how far from the sender the message will travel. The strength starts at a predefined level and degrades as it travels away from the sending cell. A signal with a lower strength will have a smaller effect on a cell.

Each cell keeps track of the signals it has received, tracking I'M DYING signals separately from PLEASE DIE signals. Each time a cell receives a signal it will remember the strength of that signal. Once the sum of those strengths passes a threshold value, the cell will undergo apoptosis. This threshold is defined separately for I'M DYING and PLEASE DIE, and represents when a cell has been “convinced” to die from the signaling. We therefore guarantee that more than one signal is required to be received before a cell undergoes apoptosis.

The flowchart including the communication can be seen in Figure 3.6. Both cancer and normal cells check if the sum of the strengths of received signals crosses the apoptosis threshold. If it has, then they may send the I'M DYING signal and then undergo apoptosis. Otherwise, they follow the normal flow of decisions, with normal cells also checking if they should send PLEASE DIE. If a cell sends a PLEASE DIE or I'M DYING message it is sent to the environment. At the end of every tick, after all cells have updated, the environment sends all signals that were sent during that time tick.

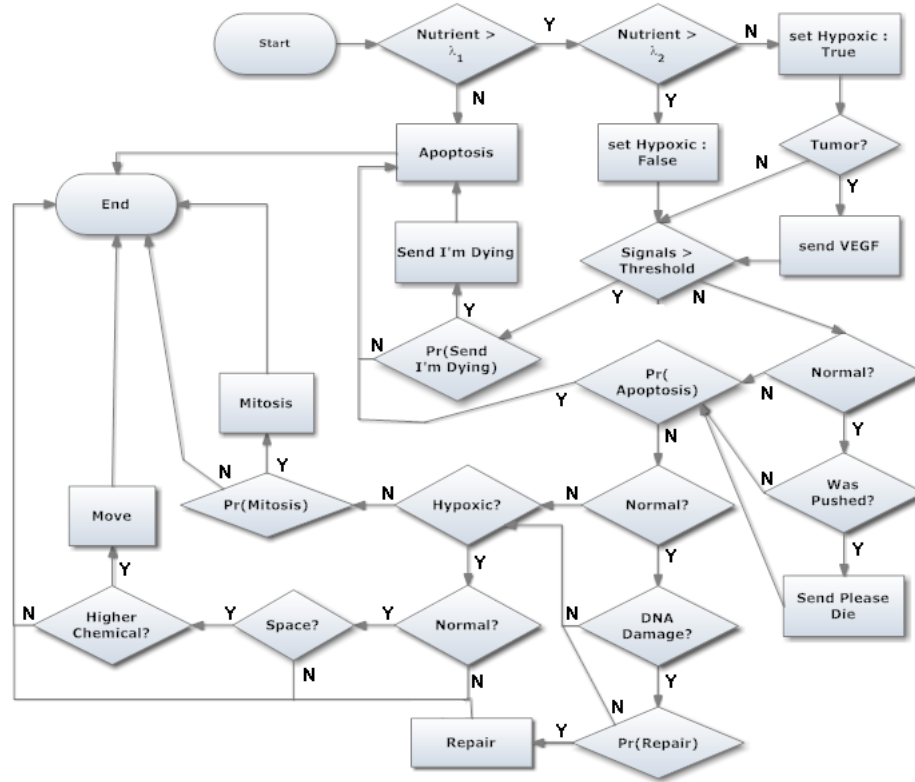


Figure 3.6. Flow chart for both normal and cancer cells including communication of I'M DYING and PLEASE DIE signals, for every time step. The differences for normal and cancer cells are what probabilities are used and how the genes affect decisions (not shown). Receipt of I'M DYING and PLEASE DIE signals is not shown as that happens in a separate process between ticks.

This mechanism keeps messages from being lost due to the cell updating being done linearly, and preserves time order.

We do not assume that cancer cells always send the I'M DYING message before they die from signaling, nor do we assume that they receive or acknowledge every I'M DYING or PLEASE DIE message sent to them. Within the model we represent this uncertainty by a probability for each failure. We also do not assume that the communication protocols are always functioning. Instead we allow the communication protocols to only activate after a certain number of cancer cells exist within the system. This type of system adversity mimics a delay in acknowledge abnormality within the area. In the Results we analyze how often these failures can occur without impacting the system's ability to remove cancer cells.

Ratio	Healthy Proliferation	Healthy Apoptosis	Cancer Proliferation	Cancer Apoptosis
6	0.05-0.15	0.0001 - 0.0002	0.3 - 0.9	0.0001 - 0.0002
10	0.03 - 0.09	0.0001 - 0.0002	0.3 - 0.9	0.0001 - 0.0002
20	0.015 - 0.045	0.0001 - 0.0002	0.3 - 0.9	0.0001 - 0.0002

Table 3.1. Cellular basic protocol parameter values used in this study. Each line will be referred to in the text by the ratio, which is the cancer proliferation rate divided by the healthy proliferation rate. Each cell may vary by up to 10% from its creator, but must remain within the given ranges.

3.5 Simulation Details

Three techniques are used to provide biological plausibility and applicability to the simulation results. They are population heterogeneity, a relative proliferation ratio, and varied experimentation. To speed up and ease the analysis of growth after the initial tumor cell develops, we plant a single cancer cell into the model at a set time for every experiment. Each set of parameters is run eighteen times, each time with a different set of random number generator seeds. The results are then averaged together to give a percent of experiments that resulted in success. Success is defined as removing all cancer cells without also removing all normal cells.

Population Heterogeneity: Variances are created by ranges of death and proliferation probabilities within each cell population, which vary by +/- 200% and +/- 300% respectively (Table 3.5). A newly created daughter cell inherits these probabilities with random skewing from the parent's characteristics. To assure probabilistic cover for feasible ranges of variables due to the difficulty of estimation, we include large ranges of values.

Relative proliferation ratio: This factor describes the quotient of the valid range of tumor and normal proliferation probabilities. For demonstrations we use ratios of 6, 10, and 20. This ratio characterizes the tumor cells in terms of how much their proliferation is increased over the normal cells. For each ratio the tumor proliferation rate is taken to be in the range of 0.3 to 0.9, and the tumor death rate is 0.0001 to 0.0002. The normal proliferation rates are in the range of 0.05 to 0.15 for ratio 6, in 0.03 to 0.09 for ratio 10, and in 0.015 to 0.045 for ratio 20. The normal death rates are within the range of 0.0001 to 0.0002 for all ratios. These ranges are based on biological data [25].

Varied Experimentation: To ensure more accurate results, each set of parameters is tested multiple times since even for a fixed set of parameters describing the system the end result may

Pairs of Thresholds					
Normal	4	4	5	6	6
Cancer	4	6	6	6	16

Table 3.2. Apoptosis threshold tested with each parameter set.

Communication Parameters						
Num	Cancer I'M DYING		Healthy I'M DYING		Healthy PLEASE DIE	
	Distance	Strength	Distance	Strength	Distance	Strength
1	1	1	1	1	1	1
2	1	2	1	2	1	2
3	1	3	1	3	1	3
4	2	3	1	2	1	2
5	2	3	2	3	2	3

Table 3.3. Parameter sets for the communication protocol.

differ as life protocols of individual cells are defined in a stochastic manner. This increases the realism of results, as it ensures that we test many different possible outcomes from the same set of basic parameters in the system. The variations on the communication protocol can be seen in Table 3.3 and Table 3.2.

We therefore experiment using these parameter sets, the basic cellular model, angiogenesis, and the communication model to determine in which cases we are able to remove cancer cells without also removing too many normal tissue cells.

3.6 Results

Through initial experiments we saw that using only I'M DYING or only PLEASE DIE does not result in removal of all cancer cells in the system. We therefore experiment with both signals being used in combination, and test the system when cancer cells fail.

As described previously, our primary claim is that *inter-cellular messaging among cells based on neighbor death and spatial impingement can be used to encourage death of surrounding cells such that primarily cancer cells are killed and healthy tissue cells survive* (Claim 2). We investigate a series of hypotheses to determine to what extent this claim is supported:

Hypothesis 2.1 *Communication parameters indicating a higher distance traveled and a stronger signal will succeed in removing cancer cells most frequently.*

Hypothesis 2.2 *Higher thresholds will correspond to lower success rates, due to increased difficulty in removing cancer cells.*

Hypothesis 2.3 *Higher proliferation ratios will succeed in removing cancer cells less frequently than lower proliferation ratios.*

Hypothesis 2.4 *There exists an amount of delay in the start of communication that will decrease the success rate, for all parameter combinations.*

Hypothesis 2.5 *Communication will succeed less frequently when cancer cells are allowed to either ignore received signals or refuse to send I'M DYING, and even less when both failures are allowed.*

We investigate these hypotheses in five scenarios: as initially described, with a delay in the start of messaging, with cancer cells ignoring received messages by some probability, with cancer cells failing to send I'M DYING with some probability, and with cancer cells both ignoring received messages and failing to send I'M DYING.

3.6.1 Communication Protocol Never Fails

First we assume that the communication protocols always function correctly. As can be seen in Figure 3.7, the first four parameter sets always succeed in removing the cancer cells. This is true across all apoptosis thresholds and all proliferation ratios. The fifth parameter set always fails, however, and will thus not be further tested. This parameter set represents the highest combination of signal traveling distance and signal strength, showing that too high of a signal can be detrimental to the system.




Figure 3.7. Rate of success in removing cancer cells. All proliferation ratios and apoptosis thresholds generate the same results.




Figure 3.8. Rate of success in removing cancer cells when there is a delay in the start of the communication protocol. The legend shows the parameter set. All apoptosis thresholds are combined, because there is no variation among them.

Since all lower distance and strength combinations result in complete removal of the cancer cells, it is possible that the choice of parameter set is not important. This finding directly refutes Hypothesis 2.1 with regard to a system in which cancer cells always follow the communication rules. Since all thresholds and proliferation ratios also give identical results, Hypothesis 2.2 and Hypothesis 2.3 are also refuted with regard to a perfectly functioning communication system. Hypothesis 2.4 is not relevant.

3.6.2 Delay of Communication Protocol Start

Next we analyze our hypotheses against a system in which the start of the communication protocol is delayed until a certain number of cancer cells are present (Figure 3.8). If this communication was activated by therapy, a delay would occur as therapy will not occur as soon as the first cancer cell develops. A delay is also possible if the cancer has a mechanism to initially block these messages.

For both proliferation ratios 6 and 10, the system succeeds in removing all cancer cells up to a delay of 12000 cancer cells (about 20% of possible total of cancer cells in system). For a proliferation ratio of 20, the system only succeeds up to a delay of about 4000 cancer cells (about 6% of possible size). Thus, Hypothesis 2.3 is supported in the case of a delay.

In all three proliferation ratios, when there is a delay we still do not see a difference among the apoptosis thresholds. This is the same as in the initial results, again refuting Hypothesis 2.2. These results do start to show a difference among parameter sets, however. For all three proliferation ratios, the highest parameter set (4) has a higher percent of success in at least one instance. For the proliferation ratio of 20, it has higher rates of success for almost all levels of delay. The parameter




Figure 3.9. For each proliferation ratio, the highest rate of failure to receive messages that still results in 100 percent removal of cancer cells (diamond), and the lowest rate of failure to receive messages that results in no removal of cancer cells (circle). For all ratios, the fourth (strongest) parameter set can succeed when tumors do not receive messages up to about 77% of the time.

set 4 is also the only instance in which cancer cells have a stronger signal that travels further than the signals sent by normal cells. In at least the case of a delayed start of the communication protocol, a stronger I'M DYING signal from cancer cells is most beneficial, partially supporting Hypothesis 2.1. At first this seems unlikely, as we do not assume cancer cells are willingly helpful in their demise. However, if cancer cells retain the ability to send I'M DYING messages but it is amplified as part of the mutation to its controls, then we could have the situation represented by parameter set 4. It is comforting that if cancer cells are able to send stronger apoptosis signals, and delay the start of message passing, they will kill more cancer cells than surrounding tissue cells.

3.6.3 Failure to Receive Messages

As discussed in the Model section, cancer cells may ignore some percentage of messages they receive. It will not add that message to its sum of strengths, and thus it will not move closer to dying from the communication protocol. We represent this chance of failure as a probability, representing what percentage of messages will be ignored.

The success rate in removing cancer cells is decreased as the rate of cancer cells failing to receive messages increases, supporting Hypothesis 2.5 (Figure 3.9). The failure to receive messages is best dealt with when cancer cells send stronger I'M DYING messages, since the fourth parameter set is the most successful for all three proliferation ratios. The success rate of the fourth and strongest parameter set supports Hypothesis 2.1 in the case of a failure to receive messages. The communication protocol can still succeed with failure rates up to 20-30% with the other parameter sets,




Figure 3.10. For each proliferation ratio, the highest rate of failure to send messages that still results in 100 percent removal of cancer cells (diamond), and the lowest rate of failure to send messages that results in no removal of cancer cells (circle). For all ratios, the fourth (strongest) parameter set can succeed when tumors do not send I'M DYING up to about 90% of the time.

and therefore we do not need to rely on cancer cells sending stronger signals to be able to remove them if they fail to receive all messages. As was shown with a delay in communication start, the thresholds do not play a role in determining success (refuting Hypothesis 2.2), and the success rate decreases slightly as we increase proliferation ratio (supporting Hypothesis 2.3).

3.6.4 Failure to Send I'M DYING

Cancer cells may also fail to follow the communication protocol by not always sending I'M DYING before dying from the communication. Supporting Hypothesis 2.5, the success rate of removing cancer cells decreases as we increase the probability of them not sending I'M DYING (Figure 3.10). They can fail to send I'M DYING up to 50% of the time with parameter sets 2-3. Parameter set 1 succeeds 100% of the time up to a failure rate of 30%. With parameter set 4 they can fail to send I'M DYING up to 90% of the time and still remove the cancer cells, across all proliferation ratios (supporting Hypothesis 2.3). This success is significant, as one would expect cancer cells to fail to adhere to the protocol most of the time.

Thresholds continue to not play a role in this scenario, again refuting Hypothesis 2.2. The strongest parameter set does show a better success rate in cancer removal, supporting Hypothesis 2.1. Since the communication does not succeed 100% in removing cancer cells when there is a 100% failure rate of sending I'M DYING, these results also show that a combination of PLEASE DIE and I'M DYING is indeed necessary to remove cancer cells.




Figure 3.11. For each proliferation ratio, the highest rate of both failures that still results in 100 percent removal of cancer cells (diamond), and the lowest rate of both failures that results in no removal of cancer cells (circle). For all ratios, the fourth (strongest) parameter set can succeed when tumors fail to adhere to the messaging rules around 60% of the time.

3.6.5 Failure to Receive Messages and Send I'M DYING

Cancer cells may also fail both by not receiving all messages sent to them and by not always sending an I'M DYING message before dying from the communication. We test this possibility by allowing both failures to occur with the same probability. For all proliferation ratios, parameter set 4 significantly outperforms the other three parameter sets (Figure 3.11), supporting Hypothesis 2.1. For all proliferation ratios it succeeds when both protocols can fail approximately 60% of the time. A success rate of 100% with 60% failure to adhere to the messaging rules is outstanding.

The other parameter sets vary between the different proliferation ratios, but overall we see 100% success from all when a combined failure rate of up to 20% is used. For the individual failures, 100% success was obtained up to about a 30% failure rate. Thus, Hypothesis 2.5 is supported as the combination of failure types decreases the rate of cancer removal.

3.6.6 Result Summary

Overall, our initial claim that inter-cellular communication techniques could explain the removal of cancer cells is supported by our computational experiments. We also tested in what scenarios and to what extent this claim is supported by analyzing the results in terms of five hypotheses. We found that the communication protocols are robust to failures. The best success is obtained with parameter set four, where only a minority of sent messages must be acknowledged by cells and only a minority of cancer cells must adhere to sending the I'M DYING message. We still see success with the other parameter sets as well, although a majority of adherence to the communication protocols

is required. The system can continue to remove cancer cells even when multiple adversities occur, which is beneficial when dealing with cancer. In cancer there are multiple failures accumulating over time, and errors in these communication protocols could easily be one of them.

3.6.6.1 Hypothesis 2.1

During initial experiments with a fully functioning system, the four weakest parameter sets show identical success rates of 100%, whereas the strongest parameter set never succeeds. However, as we perturb the system with delays, failures to receive messages, and failures to send messages, we find that the strongest tested parameter set (the second strongest in the original set of five) succeeds more frequently than the weaker parameter sets. Thus, it is the case that an increase in strength of signal does not affect a correctly functioning system, but can improve the performance in a system with faulty message passing. In both cases though, if a signal is strengthened too much it is detrimental instead of helpful.

3.6.6.2 Hypothesis 2.2

Throughout all scenarios tested, there was no significant difference between the results between the different thresholds. It is surprising that we do not see a difference, as a higher apoptosis threshold makes it more difficult to kill cancer cells, and should thus decrease the removal success rate. This difference is possibly not occurring because there are enough signals in the system, even with a delay or failure to adhere to the communication protocol, that cancer cells are well enough saturated in them. Despite that, there are not enough signals in the system to kill too many healthy normal cells in addition to the cancer cells.

3.6.6.3 Hypothesis 2.3

A higher proliferation ratio represents a more aggressive cancer. An increase in proliferation ratio decreased the success rate only when the communication was perturbed by a delay in activation, failure to receive messages, or failure to send I'M DYING messages. With a fully functioning communication system there was no difference in the results for different proliferation ratios.

3.6.6.4 Hypothesis 2.4

A delay to the activation of the communication rules above a specific threshold results in a decrease in the success rate of cancer removal for all parameter combinations. We find that for both a proliferation ratio of 6 and 10, the system succeeds in removing all cancer cells up to a delay of at least 12000 cancer cells (about 20% of possible total of cancer cells in system). For a proliferation ratio of 20, the system only succeeds up to a delay of about 4000 cancer cells (about 6% of possible size).

3.6.6.5 Hypothesis 2.5

There is no evidence to refute this hypothesis. Whether cancer cells are allowed to ignore received messages, not send I'M DYING messages, or a combination of both, these types of failures to the communication rules result in a decreased rate of success in removing cancer cells.

3.7 Conclusions

We present a new three-dimensional agent-based model for cancer growth. This model is an improvement over previous models as it explicitly models normal tissue cells in addition to cancer cells. The functioning of cells is probabilistic, with probabilities drawn from the literature and experts in the field. Additionally, the definition of the functioning of a cancer cell is based on the hallmarks of cancer [65]. We directly model nutrient diffusion from vasculatures, and cancer angiogenesis through the diffusion of VEGF that initiates sprouts from blood vessels.

We also propose a pair of local communication protocols for removal of cancer cells. Both types of messages are remembered by receiving cells, and once enough are received the cell will undergo apoptosis through a secondary mechanism. The first message, I'M DYING, is sent just before the cell undergoes apoptosis due to the communication, and is inspired by phenomena seen in chemotherapy. The second message PLEASE DIE provides a mechanism for activation of the I'M DYING message. It is sent only by normal tissue cells when they detect that another cell has pushed them to the side, either due to movement or proliferation.

This combination of message types is tested to determine how well they work together to remove cancer cells from the system. Five variations on strength and distance for signals is tested against three variations on cellular proliferation ratios between cancer and normal cells, representing vari-

ations in the aggressiveness of cancer growth. Initial tests show that without any perturbations to the system, all variations but one on strengths and distance for the signals results in 100% success in removing all cancer cells in the system without decimating the normal cell population. The one parameter set that does not succeed is the strongest parameter set, showing that the strength and distance of the messages must be limited.

We also test perturbations on the messaging system. The first perturbation is a delay on the initiation of the communication system, based on the number of cancer cells in the system. Results show that the system can handle about 12000 cancer cells before activation for the lower two proliferation ratios, and 4000 cancer cells for the highest proliferation ratio. These results indicate a window of opportunity for re-activation of inactive communication protocols in our model in which they can still successfully eradicate the tumors.

We do not assume that cancer cells always adhere to the communication protocol. One way we allow them to fail is by ignoring messages sent to them by other cells, with some probability. The strongest parameter set tested in this scenario can handle a cancer communication failure of up to 70%. We allow cancer cells to fail by not sending I'M DYING as well. The same parameter set can handle cancer cell failure to send I'M DYING up to 90% of the time, and still remove all cancer cells from the system. When we combine these failures the results vary by proliferation rate. For a proliferation rate of 6 the system can handle 60% failure rates. The higher proliferation ratios can only handle 40% failure rates. Therefore, we do not need to assume that cancer cells will always follow the communication rules, but only that they follow them at least some part of the time. How often they must follow the rules depends on the aggressiveness of the cancer only when both types of failures can occur.

We test a variety of parameter values for the communication protocol to test a variety of possibilities within biology. We do not know which rates of message sending are most accurate, but we see that all but one of our tested parameter sets is able to remove the tumors even with perturbations to the system. Overall we see that the parameter set four is very robust to failures in the communication protocol. This parameter set represents when cancer cells send a stronger message that travels further than in the other parameter sets. It is not unlikely that a cancer cell could send a stronger apoptosis inducing signal than a normal tissue cell.

The role of the tumor microenvironment in cancer development and progression as well as in clinical outcome is increasingly acknowledged, and can now be viewed at the genomic level (summarized in [47]). For the tumor-environment interaction suggested in our model, there are emerging data to support the significance of cancer cells reporting their death. In vitro models further supported by clinical outcome data showed that HMBOX1 release by tumor cells exposed to chemotherapy is essential in mediating effective tumor eradication by different chemotherapy agents and solid tumor models; HMBOX1 release functions to report initial cell damage by chemotherapy to the immune system and likely cooperates with additional messengers released by dying cells as its artificial introduction does not lead to the same effect [9, 121]. The computerized I'M DYING signal likely parallels this biological mechanism.

We hypothesize that these communication mechanisms may already exist in nature, but the onset of cancer is due to a breakdown of the messaging either in ways outlined in our results or others. If this type of communication is found to represent communication occurring within cells, then the next question would be to determine what causes them to break down and allow cancer to grow, followed by a question of how we can re-introduce the signals as a new therapy. We have shown that the signals could succeed in removing cancer cells in a variety of situations, and that they can be reactivated later in the tumor's growth (delay results).

CHAPTER 4

MULTI-AGENT FAULT TOLERANCE INSPIRED BY CANCER

4.1 Introduction

We investigate fault tolerance for a system of cooperating agents. For a multi-agent system to function continuously it must adapt on-line to failures. There are essentially three different ways in which a system can fail: unreliable infrastructure (system supporting the agents fails), non-compliant agents (agents ignore protocols), and emergent dysfunctions (small program errors lead to real failures). Although each of these types of failures may cause different problems for the system, in each case the same high level process needs to be followed to deal with the problem: detect and diagnose the problem, and then fix the problem. We focus on fixing the failure of non-compliant agents.

The two main approaches for fixing these three types of failures in multi-agent systems are survivalist and citizen. The survivalist approach requires each agent to be capable of dealing with all problems as an individual agents that follows a prepared set of actions for each specific problem [90]. For instance, a set of replicas could be maintained and then deployed once a fault is detected. However, this approach requires the designer to be able to anticipate all types of faults within the system and cannot easily deal with agent death. The citizen approach is the other extreme, as it utilizes an external system that is alerted when an agent dies and then reallocates tasks so that the overall system continues to function correctly [80]. Thus, the citizen approach can very easily deal with agent death. However, there is now an additional system that must be maintained, which creates a single point of failure.

My approach is a combination of these two techniques as a Citizen Group approach. It differs from the survivalist approach as it does not require all agents to deal with failures individually. Unlike the citizen approach that requires special monitor nodes to diagnose failures, my system has each agent monitor its neighbors to detect and eliminate anomalies. This citizen group approach is accomplished by defining a cancer-inspired mechanism for multi-agent systems that improves

robustness by enabling agents to combat malfunctioning agents. In the cancer model we saw that local anonymous broadcast communication could be utilized to remove cancer cells without necessarily knowing which cells were cancerous. Based on the fact that cancer cells grow in a cluster this mechanism removed all cancer agents by only detecting the irregularities of a few of them. Seeing that this mechanism works well to remove cancer, we apply the same techniques to the problem of multi-agent fault tolerance. In this case, we examine agents instead of cells, and determine how to remove malfunctioning agents instead of cancer cells. This communication mechanism is described in terms of a system we call HADES, “Healing and Agent Death Enabling Stability.”

HADES includes self-regeneration through the creation of new agents. In biology, this is called proliferation and is used to both recover from dying cells as well as to develop and expand the system [95]. In computational systems regeneration is used in two different ways: the initial creation of replicas that function throughout the life of the system and can replace failed agent when necessary, or cloning (creation of a copy) of the agent when another agent is needed. In [49] agents replicate to create agent clusters that make joint decisions, thus improving fault tolerance. For hardware, a simple robot composed of building blocks can generate a new copy of itself by assembling additional blocks [157]. In multi-agent systems, cloning is often used to create a copy of the agent when an agent is overloaded, so that it can share its load with its clone [76, 136]. In HADES, we are utilizing regeneration in the form of cloning, where a new agent is only created when it is needed, and is created as a direct copy of the cloned agent. Additionally, it is the cloning agent itself that decides when it will be necessary to clone itself.

In HADES all agents follow the same basic protocols that control individual activity of the agents and provide generalization and improvement over previous self-regenerative agent systems. In addition, HADES diagnoses and repairs via a multi-step protocol. The first step is self-monitoring such that an agent determines whether its own basic protocols have been damaged. The second step is for the agent to repair any discovered damage within itself. It is possible that the agent will be unable to repair itself and will thus apply self-death to retain the system’s health, so that its failures will not harm the system goals. The fourth step comprises a main concern of this work: the ability of agents to note that at least one of their neighbors is irregular and thus send surrounding agents a message as a warning. These messages cause the receiving agents to either elevate their own level of alertness or if enough messages are received they entice the receiving agents to activate their

programmed death. As will be described in the paper, all agents maintain some level of citizenship and communicate their upcoming death via signaling to entice neighboring agents to die as well. HADES studies the ability to induce self-death on an agent, allowing the system to recommend death without agents killing one another. It utilizes agent regeneration, repair, and death, as well as novel communication protocols to overcome system faults and maintain continuous operation. All the activity of HADES is cheap and will only occur as necessary, and thus can serve as an underlying maintenance mechanism to improve the robustness of any multi-agent system.

This chapter is organized into the following sections: Section 2 describes more related work in diagnosis and repairing systems, Section 3 introduces HADES as a multi-agent system where agents follow goals and life protocols, Section 4.4 describes how an agent may become irregular and how its continued operation may negatively impact the entire system and includes the communication protocol that enables the recognition and removal of such irregular agents by their neighbors, and results are described in Section 4.6. We close with conclusions.

4.2 Related Work

For a multi-agent system to function continuously it must adapt on-line to changes in the environment and to internal failures. Diagnosis of a problem is a key requirement, as is providing a plan to react to the problem [64]. Various frameworks exist for diagnosis in multi-agent systems, including domain independent diagnosis where an agent should also be able to determine a new plan if its expectations are not met [72]. Diagnosis for pre- and post-failure analysis for causal tasks can allow the system to both prevent a failure and recover from it [150]. It is argued that post-failure protocols are less domain dependent and are more crucial for the design of robust systems [150].

One way to obtain post-failure robustness is via *self-repairing* mechanisms. Repair typically follows one of two categories, “Attributive” or “Functional.” Attributive repair restores attributes to their state prior to the failure to revert any damage. Functional repair does not backtrack but instead optimally uses the remaining resources to obtain the best possible functionality [33]. As an example of Attributive repair, software components can self-monitor to determine vulnerabilities and thus remove them [131]. The analysis to determine vulnerabilities uses both static and dynamic techniques, including the Stackguard tool and predicate abstraction. Software validation techniques can be used to identify causes of vulnerabilities, enabling their removal [131]. Wireless sensor

networks are also being designed with self-healing capabilities inspired by immunology to detect sensor faults and respond via a form of Functional repair. For example, [16] mimics B-cells in the immune system with scripts on monitor nodes that follow the status of the sensor nodes in the system. The monitor nodes can find failure by examining the statistical properties of the sensor readings. This system can adapt to the changes in the network caused by sensor failure via monitor nodes notifying sensors of incorrect readings, allowing them to request retraining. This combination of different node types interacting enables the system to find and react to failures [16].

Self-regeneration provides another paradigm for attaining robustness, and has been investigated for at least the last 50 years [152]. It works in a functional framework, mainly to achieve a larger system during development or after agent death. It is one of the main responses to agent death utilized in other systems. Roth et. al. introduces a system that uses self-regeneration to self-organize, growing from a stem cell into an intelligent behaving organism [126]. The system is comprised of cells that replicate based on chemical signaling received from the surrounding environment. Since the replication control is sensitive to the environment it will cause the system to grow from a single stem cell to the desired size. This algorithmic control will also cause the system to regenerate after cells are (artificially) removed. Replicating upon need is considered a system-level repair mechanism and is different from repair mechanisms of the cell itself, which we include in HADES. The focus in [126] is on the growth and organization of the system as opposed to the later fault processing that is the main contribution of HADES.

The use of self-regeneration for development is described in [97]. In this system artificial chemicals are used to control movement of cells on a lattice, and grouped with the regenerative abilities the system is able to retain a specific shape despite dying cells. The ability of the system to obtain a specified shape despite dying cells is referred to in that work as self-repair. This use of the term self-repair is to be differentiated from the more fundamental definition, as is used in biology and will be used in our system: the cellular level repair. In cellular level repair, a self-replicating cell or agent uses error correction to make changes to itself to ensure that it grows into the proper form [61]. The system self-repair is but one consequence of the agent-level repair.

Self-regeneration can also replace dying agents in a general three-dimensional lattice using artificial presence chemicals [56]. When an agent dies, the lack of chemical in its location will cause the neighbors to regenerate to solve the problem of agent death. In HADES the use of presence

chemicals is advanced to a more efficient form by allowing the presence chemicals to linger rather than removing them with each step. This enables our agents to also recognize the center of the shape and thus keep a coherent shape without using the previously required ordering [56]. HADES includes a deeper sense of self-death, mutations, and communication protocols.

In [88] regeneration, repair, and death are combined to create an artificial organism as the first step toward hardware with the ability to remove surrounding agents that may be faulty. The cells in the system are organized in a grid. One form of repair involves disabling all cells in the column of the faulty cell after transferring their functions to the cells in the column to their right, essentially using death to repair the organism. The other form of repair is internal, used to combat failure of the artificial molecules that control cellular actions. This repair is accomplished by removing the faulty molecule and then rearranging the remaining ones until a spare is reached, ending with the same number of molecules as was used before the failure [88]. This approach lacks sufficient flexibility in agent movement. It also requires spare agents to be kept to the side until needed, instead of creating them as necessary. This system does not give cells the ability to decide their own death, which we know from biology and multi-agent systems can be quite dangerous. We thus improve by utilizing similar mechanisms overall, such as internal repair, regeneration, and death, but with allowing the agents to decide when to die themselves. The fact that they keep spare agents to the side is something for us to consider should we move from simulation to real system, although from [144] we have seen that there are other ways to accomplish the same effect.

Agent death is not only a problem to the system, but can be a desired property when the system has to decrease in size or when agents are destroyed and act irregularly. When death is desired, self-destruction is preferred over agents killing other agents, providing more robust behavior [80]. [144] describes a self-managing system proposed by NASA that uses self-destruction as a last resort to deal with damage. This system first tries to self-repair, but if the repair fails a self-death mechanism is deployed to remove the broken agent. This self-death is implemented by a constant *stay alive* signal, such that if an agent no longer receives the signal it will self-destruct. In [88], regeneration, repair, and death are combined to create an artificial organism as the first step toward hardware with the ability to remove surrounding agents that may be faulty. One form of repair involves disabling all cells in the column of the faulty cell after transferring their functions to the cells in the column to their right, essentially using death to repair the organism. The other form of repair is internal, used

to combat failure of the artificial molecules that control cellular actions. This repair is accomplished by removing the faulty molecule and then rearranging the remaining ones until a spare is reached, ending with the same number of molecules as was used before the failure [88].

HADES entices death as well but via the use of communication protocols, thus not requiring agents to know exactly which other agents are faulty. In [144] self-death is a default that occurs if an override is not received. Since irregularity is not the norm in most systems, this technique will continuously flood the system with messages to every agent. However in our system messages are utilized in the opposite way by being sent when irregularity has occurred and only to agents in that area, thus decreasing the overall messages that must be sent. We are thus able to maintain our system by suggesting death to surrounding agents and over time removing all faulty ones. Overall, HADES uses a unique combination of death, repair, regeneration, movement, and communication to retain itself by giving agents the utmost control over these actions.

4.3 HADES - The System

HADES is a three-dimensional lattice of agents that act locally to achieve citizenship goals. Agents clone themselves to replace any neighboring agents who have been removed from the system to ensure a complete system as often as possible. To decrease the number of times agents clone themselves, and thus reduce the chance of error over time that occur due to the cloning process, agents do not clone themselves immediately when a neighbor disappears.

Agents also aim to maintain self-health by testing for failures that are represented by mutations to life protocols and either repairing them or inducing self-death if they are not repaired. Each agent also has a goal to retain basic distance from surrounding agents. This distance is maintained by not moving or replicating into a location closer to its neighbors than the specified distance. The system also aims to maintain itself as one cohesive unit. Agents move or create a copy of themselves in a direction closest to the highest density of surrounding agents, to ensure that agents do not become isolated. Finally, the agents also initiate and transfer signals among themselves to facilitate the goal of maintaining the overall system health by decreasing the number of irregular agents.

To achieve these goals each agent follows a series of basic protocols that focus on the internal state of the agent. These protocols are cloning, self-testing, repair, suppression, self-death, space maintenance, and movement. The basic protocols are highly inter-regulated, as will be seen below.

Other agent protocols consider the environment and communication with neighbors, and will be described in section 4.4.1. The basic protocols are inspired by cellular biology, but are applicable for other self-regenerating multi-agent systems as each protocol relates to a need within multi-agent systems. We next describe how they are used to achieve an agent's goals, and how they each represent general multi-agent systems.

4.3.1 Agent Cloning

All agents share the goal of keeping the equilibrium of the system. This goal is accomplished by cloning, which relates to the biological mechanism of proliferating cells. A probability controls how frequently each agent attempts to clone itself if there is a neighboring unoccupied location, to ensure that it neither clones too quickly nor too slowly. The rate of cloning is a system parameter that can be changed and the Results section will show how it affects the state of the system. Mutation to a single basic protocol may occur with small probability during cloning, which will cause the newly created agent to function incorrectly in some way. This alteration to the agent's functioning is further discussed in Section 4.4.

Since cloning is a fundamental operation and its failure can cause significant damage, there is a specialized mechanism called the "suppression" protocol for controlling the effect of mutating the cloning protocol. If an agent attempts to clone itself more frequently than the given probability allows, the suppression mechanism will halt the replication. Thus an agent can only create too many clones of itself if both the cloning protocol and the suppression protocol are mutated, thus utilizing double mechanisms for agent robustness.

4.3.2 Internal Agent Repair

Agents maintain their own functioning by monitoring any damage to the protocols from mutation during cloning. When a mutation is detected internally, the agent attempts to repair the damaged protocol. If the repair mechanism continuously fails, the agent recognizes that it may not be functioning correctly and will attempt to remove itself from the system so that it does not damage the system. The agent is therefore preserving the system health by preserving its own health. In the case where the death protocol is damaged the use of external communication will be required, as will be discussed later in the chapter.

4.3.3 Maintaining Enough Agents

Each agent in HADES requires a spatial buffer around itself. No normal agent will move to a coordinate if another agent exists on an adjacent location. Also, a daughter agent will only be created if there is an empty location with no agents adjacent to it. If an agent becomes adjacent to another agent despite these efforts, that agent will attempt to move away if it can do so without encroaching on another agent's space.

Space is not usually the deciding factor in an agent system. However, this spatial buffer is an analogy for multiple aspects of a multi-agent system: the requirement of a certain amount of agents to perform the needed tasks, and for agents to be properly dispersed among different aspects of the task to properly run in a distributed fashion. Thus, if there are too many agents within an area, then there are a lot of wasted resources and the agents that need to perform that area's duties will not be able to succeed. Also, if there are too few agents in the system then it will not be able to perform all tasks.

4.3.4 Presence Signals

An agent maintains the cohesiveness of the system by maintaining the shortest possible distance between itself and the center of the system. This organization is accomplished via presence signals that all agents passively emit. The areas closest to the agent will therefore have a stronger presence. Presence signals are continuously diffusing in all directions for a specific radius and time period. These signals can thus be used to determine the proximity between agents, and the direction of the center of mass.

Presence signals will diffuse as in Algorithm 4.1, with the concentration of the chemical slowly increasing at each time step by $\frac{\epsilon}{MAX_RADIUS^\rho}$ for each location within the distance. We used a *MAX_RADIUS* of 5, ϵ of 1, and ρ of 2. *MAX_RADIUS* represents the total radius a presence signal can travel, and thus once the signal reaches this distance from the sender the signal strength will be steady at each location until the emitting agent moves or dies. When an agent moves or dies, the signal will slowly decrease toward the original spot at the same rate that it diffused out.

The presence signals act with cloning and spatial availability to maintain the correct number of agents within the desired organization of space. Since agents are allowed to move to a neighboring

Algorithm 4.1 Presence Signal Diffusion

```
1:  $S = time_{current} - time_{lastmove}$ 
2: if  $S < MAX\_RADIUS$  then
3:    $distance = 0$ 
4:   for  $S > 0$  do
5:     for all location(x,y,z) within distance (city block distance) do
6:        $Chemicals(x, y, z) = Chemicals(x, y, z) + \frac{\epsilon}{MAX\_RADIUS^p}$ 
7:     end for
8:      $S = S - 1$ 
9:      $distance = distance + 1$ 
10:  end for
11: end if
```

location (or task) if they are too close to another agent, these signals help them determine where the highest concentration of other agents exists.

4.4 Irregular Agents and the Communication Protocol

In the previous section we described the basic protocols of an agent. In this section we introduce the second type of protocol, the *rescue protocol*. This protocol exists primarily to counteract agents that have obtained mutations but have neither repaired those mutations nor removed themselves from the system. Although every individual mutation can be problematic, the worst case is when all basic protocols are mutated and the agent becomes *irregular*. An *irregular* agent will:

1. Continuously attempt to clone itself, within some higher probability
2. Refuse to repair itself
3. Not remove itself when repair fails
4. Ignore spatial/resource regulations

An irregular agent will continue to clone itself since it can neither repair nor kill itself, thus spreading its damaged basic protocol to its daughters. This behavior will quickly create a cluster of irregular agents. The probability of creating an irregular agent from a normal agent is incredibly low (based on the cloning and mutation probabilities) since to prevent the agent from repairing mutations the basic protocols must be ruined in a particular order: first it must lose the ability to repair errors, then it must lose the ability to remove itself due to that repair failure, then it must lose the ability to slow down its cloning rate, and last it must lose the mechanism that prevents it

from over cloning. Once an irregular agent has been introduced into the system it will create too many clones of itself, paying no respect to available space or diffused presence signals and will quickly take over the system, similar to what was shown in the previous section with cancer cells. An irregular agent in essence “pushes” a normal agent into an adjoining lattice location if it is in its way, causing it to lose its desired space, which may hinder its ability to clone itself. Therefore, the normal agent count will slowly decrease since they cannot replenish the numbers.

An irregular agent is essentially stealing resources and tasks away from correctly functioning agents, as they grow to a relatively large number in a single area of the system and do not participate in any system goals. Agents are not able to know if their neighbor is irregular, although if a “push” occurs they will know that there is an irregular agent nearby. Since irregular agents form a cluster that pushes agents closer to their neighbors as it expands, they will continue to exert physical pressure on the same area of the system. We propose taking advantage of this style of growth to find a solution that inhibits the irregular agents and supports the normal agents, based on the communication paradigm developed for removing cancer cells in the previous chapter. The environment and signaling algorithms will therefore enable the agents to become active citizens through maintaining their own health, monitoring neighboring agent health, and maintaining the system health. This rescue protocol sends signals that encourage neighboring agents to die. The goal is to determine how these signals should be sent such that healthy agents will remain while irregular agents are removed.

4.4.1 Rescue Protocol Communication

We propose a *rescue protocol* aimed at saving the system from the over-creation of irregular agents. Agents that suspect irregularity will alert surrounding agents by sending a signal. The receiving agents interpret this signal as a statement of irregularity in the area. If an agent receives this announcement multiple times it will conclude that it may be the problem and should activate its programmed death. Agents do not know the sender of the signals they receive, and when they send a signal it is to all agents in the surrounding area. This communication style is inspired by biology [50].

The first communication protocol within the rescue protocol is called PLEASE DIE, and is initiated by healthy agents that sense irregularity when their space is invaded by a neighbor’s “push.” This type of invasion occurs when a neighboring agent moves to the space an agent is occupying, causing

that agent to be moved to surrounding buffer space. Due to normal agents maintaining space the invading agent will either be an irregular agent or a normal agent that has been pushed and therefore forced to move. If a pushed agent does not have empty space around it, it will push a neighboring agent in the same direction. This chain effect will continue to occur until an agent is pushed that has an open space surrounding it.

It is beneficial for an agent to only die after a specified number of signals (> 1) are received so that no one agent can directly kill another one. For this purpose a threshold is used:

A Threshold of Signals (TOS) for signal type S is a local variable such that:

1. an agent may only activate self-death from receiving signals once the sum of all received strengths of type S reaches the TOS for type S .
2. the TOS for type S is greater than the max strength of a signal of type S .

4.4.2 Double Messaging

One may assume that the PLEASE DIE signal suffices for restoring the system. However, this assumption is not the case as shown in Chapter 3. We thus propose a novel double message system where in addition to the PLEASE DIE signal generated by a pushed normal agent, a secondary signal is applied. This second signal called I'M DYING is sent by an agent when it is dying from receiving a signal. The I'M DYING signal is used to alert neighboring agents that they may want to consider dying as well. This method takes advantage of the cluster structure of irregular agents: neighboring agents may be descendants or ancestors of the dying irregular agent, and are therefore likely to also be irregular.

For the rescue protocol to remove all irregular agents and retain a large percentage of healthy agents the I'M DYING signal must be sent by both healthy and irregular agents that die. When no agents send I'M DYING signals the system collapses with no normal agents remaining. When the irregular agents can not send I'M DYING signals and only the normal agents send them, all normal agents are killed as well. When dying normal agents do not send an I'M DYING signal the irregular agents will eventually increase while the normal agents decrease. Therefore, it is necessary to have the PLEASE DIE signal and all agents must be able to send the I'M DYING signal for the system to be able to remove the irregular agents.

Algorithm 4.2 Signal propagation for rescue protocol

```
1:  $distance = 1$ 
2: for  $distance < radius$  do
3:   for all (x,y,z) location  $\in$  distance do
4:     send signal of strength:  $strength - (distance * \frac{1}{2*radius})$ 
5:   end for
6:    $distance = distance + 1$ 
7: end for
```

When an agent receives a signal it will always correctly identify the signal type. The two types of signals exist together, but are not combined and do not affect each other. They each have their own TOS value, as well as their own parameters.

4.4.3 Signal Parameters

Rescue signals are sent in the system for a specific radius in all directions. The strength decreases slightly with each additional distance the signal travels before reaching the radius.

The *radius* of a signal defines the maximum distance over which the signal is received. The *strength* of a signal represents the value that is associated with it, and is decreased as it travels further from the source.

Both PLEASE DIE and I'M DYING signals travel as in Algorithm 4.2. The strength and radius are varied in experiments. PLEASE DIE is only sent after a “push,” whereas I'M DYING is only sent when an agent is dying due to receiving signals. Both signals can also be received at any time step and may induce self-death. A low radius and strength coupled with a high TOS will result in many signals needing to be received before an agent will decide to die. However, a high radius and strength coupled with a low TOS value will have the opposite effect. Radius and strength can be different for each type of signal. Precise definitions of high and low for these parameters are determined in Section 4.6.

We will initially assume that all agents will die when the total of signals received reaches the appropriate TOS value. Given the mutation probability, very few irregular agents would also mutate the rescue protocol if it was an option. Also, such a mutation is only harmful if it occurs in an irregular agent, since deaths of healthy agents are considered unfavorable. However, we also test how well the rescue protocol works for removing irregular agents if irregular agents do not always follow the protocol, as well as if there is a delay in agents sending the initial signals.

4.5 Model Flow

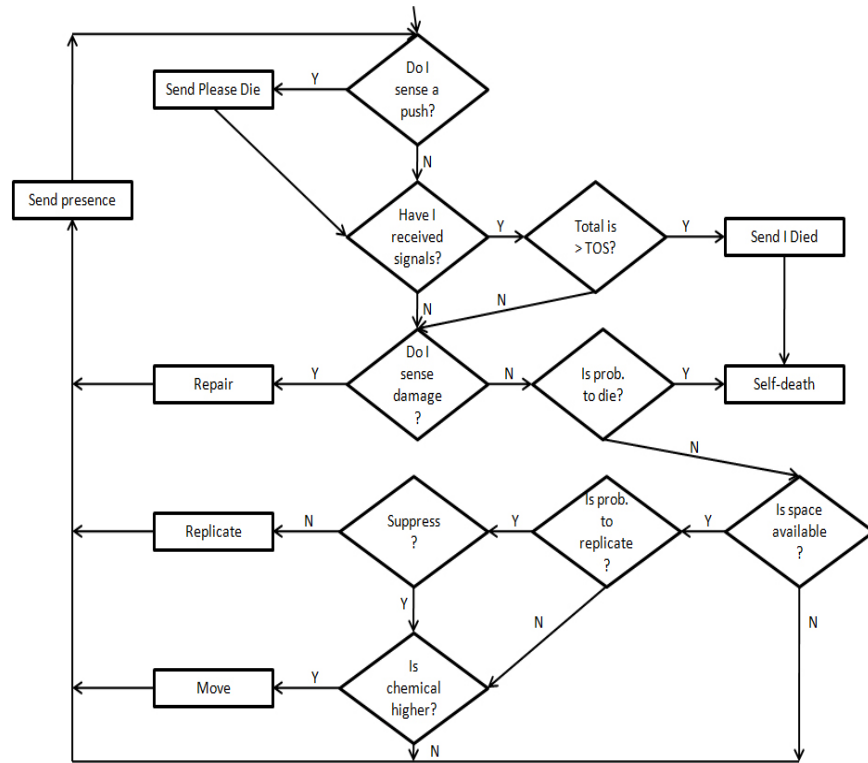


Figure 4.1. The basic action controls for a healthy agent at each time step.

The simulation progresses by time ticks, and although agent decisions are implemented sequentially, the actual actions occur in parallel. There are two different sets of actions that can be performed: at each time step one action from the first set may be chosen, and multiple actions from the second step may be chosen (see Figure 4.1). The first set of actions, only one of which may be performed at a time is:

1. Repair: occurs if the agent is damaged. With every repair attempt there is a 0.5 probability of failure.
2. Death/Removal: occurs by three mechanisms. Death can occur when an agent has been unable to repair its basic protocols. It can also occur with a small probability, to include other causes of death such as age. The third mechanism is via rescue protocol death signals sent by surrounding agents.

Ratio	Normal Cloning	Normal Death	Irregular Cloning	Irregular Death
6	0.05-0.15	0.0024 - 0.0048	0.3 - 0.9	0.0001 - 0.0002
10	0.03 - 0.09	0.0024 - 0.0048	0.3 - 0.9	0.0001 - 0.0002
20	0.015 - 0.045	0.0024 - 0.0048	0.3 - 0.9	0.0001 - 0.0002

Table 4.1. Agent parameter values used in this study. Each line will be referred to in the text by the ratio, which is the irregular cloning rate divided by the healthy cloning rate. Each agent may vary by up to 10% from its creator, but must remain within the given ranges.

Rescue Protocol Parameters						
Num	Irregular I'M DYING		Normal I'M DYING		Normal PLEASE DIE	
	Distance	Strength	Distance	Strength	Distance	Strength
1	1	1	1	1	1	1
2	1	2	1	2	1	1
3	1	2	1	2	1	2
4	1	3	1	3	1	3
5	2	3	1	2	1	2

Table 4.2. Parameter sets for the communication protocol.

3. Cloning: occurs if there is available space and suppression is not activated. The preferred probability of cloning is tested in Section 4.6. With each instance of cloning there is a 0.0001 probability of mutation. It is randomly decided with equal probability which basic protocol to mutate among all protocols that are not currently mutated.
4. Movement: occurs if an agent cannot replicate but there is an available adjacent space with a higher concentration of presence signals than its current spot, representing a space closer to the center of mass.

In addition to one of the above operations, at each time tick an agent is also able to perform the following actions in parallel: emit presence signals, and participate in rescue protocol.

4.6 Results

4.6.1 Simulation Details

For the following simulations we ran HADES as a cube bounded to size 40x40x40 units. A total of 8,000 healthy agents can reside in this cube at any given time since we chose the inherent

Pairs of Thresholds					
Normal	4	4	5	6	6
Irregular	4	6	6	6	16

Table 4.3. TOS tested with each parameter set.

distance between agents to be 2. Thus, agents only reside on every other coordinate and consider the surrounding empty spaces as buffers. However, since irregular agents do not adhere to these space constraints, they can total 64,000.

Simulations were conducted to determine the signal parameters, replication probability, and various algorithmic properties that would result in eliminating the threat of irregular agents and retaining the population of healthy agents in the given scenario. All simulations ran for 8000 time ticks. Twenty different runs were executed with different random number generator seeds for each set of parameters evaluated. All probabilities are implemented via a pseudo random number generator. We chose the Mersenne Twister RNG, as it has a long period ($2^{19937} - 1$) and “good” randomness [93]. Although the TOS values for the PLEASE DIE and I’M DYING signals are separate, they were tested with the same values. Therefore, a given TOS value refers to the value that the PLEASE DIE TOS and the I’M DYING TOS each have separately, and they are not combined in any way. In each run we recorded the number of both irregular agents and the healthy agents since both measurements are needed to accurately describe the state of the system.

The agent values used can be seen in Table 4.1. These are divided into three different replication ratios. The ratio of six denotes that irregular agents replicate six times more frequently than a normally functioning agent, which is the best case scenario. The worse case scenario of twenty allows irregular agents to replicate 20 times more frequently than normal agents. Within each ratio is a given range representing the minimum and maximum allowed, with agents varying by up to 10% from their creator’s value. Additionally, irregular agents are less likely to undergo death, but are more likely to remain in the system causing problems. This is again a worst case scenario.

4.6.2 Removal of Irregular Agents

We first analyze in what scenarios the communication is able to remove irregular agents from the system. In Figure 4.2 we see that in most cases we have a high percentage of experiments that ended in success, i.e., all irregular agents were destroyed without destroying all normal agents. The

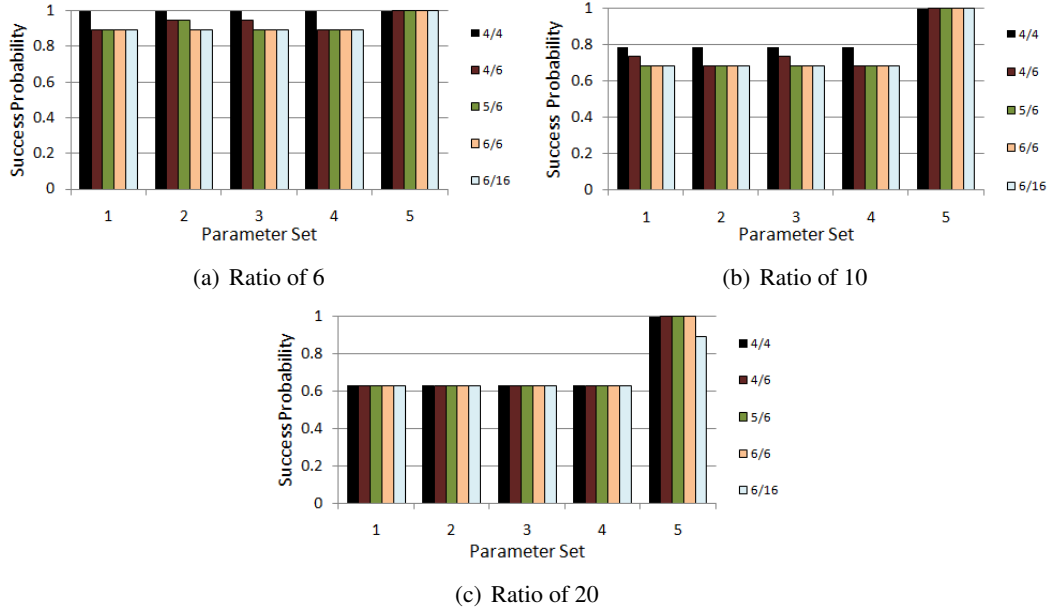


Figure 4.2. Percentage of experiments with success, defined by all irregular agents removed without all normally functioning agents removed. See Table 4.6.1 for parameter sets. Each bar represents a different threshold for dying from signaling, denoted by “normal/irregular.”

case of lower replication ratio (Fig. 4.2(a)) resulted in higher percentage of success than the middle replication ratio (Fig. 4.2(b)), which resulted in a higher percentage of success than the highest replication ratio (Fig. 4.2(c)). However, even in the highest replication ratio we see that overall there is at least 60% success in removal of irregular agents.

It is also shown that a higher TOS value does not necessarily equate to a lower success rate. In most cases there is no decrease in success when the irregular agents require more signals to die, or when both types of agents require more signals to die. We would expect a higher TOS value to correlate with lower success rates. In all but the ratio of 20 we do see a slightly decrease in success rate as we raise TOS.

The highest parameter set, correlating with the highest distance and strength for both signal types, leads to 100% success almost across the board. The only case in which this statement doesn’t hold is with the highest ratio of 20 and the highest TOS for both agent types. It therefore seems likely that this higher distance and strength is the ideal set of parameters for removing irregular agents in this type of system.




Figure 4.3. The percent of experiments with success on the most successful dataset when irregular agents either ignore signals received (blue) or do not send I'M DYING messages (green). The system can still succeed in removing all irregular agents if either irregular agents ignore at most 68% of messages received, or they do not send I'M DYING messages 72% of the time.

4.6.3 Robustness

We do not want to assume that irregular agents will always follow the rules. Thus, we also allow them to fail in two different ways: not sending I'M DYING signals when they die from the communication, and ignoring the PLEASE DIE and I'M DYING messages they receive from their neighboring agents. As can be seen in Figure 4.3, we can allow these agents to fail in each of these ways the majority of the time and still succeed in removing them. Thus the communication protocol for this system is quite robust to each agent failing to follow the protocol when they have already failed in the defined ways that make them irregular.

We also allow the protocol itself to only activate once a certain percentage of the system has been overrun by irregular agents (Figure 4.4). If we wait until 50% of the system is inhabited by irregular agents to activate the communication, we are still able to remove them completely with either the lowest or highest TOS. However, past 50% with the highest TOS values we are no longer able to guarantee success, although we can guarantee it for the lowest TOS. As we continue to delay until the communication is started, the percent of success decreases. However, even when we delay until 80% of the system is overrun, we are still able to remove the irregular agents from the system. Thus, the system is robust to the communication protocol being slow to activate.




Figure 4.4. The percent of experiments with success on the most successful dataset when signals are not used until some percentage of the system has been overrun with irregular agents. Even when waiting until 50% of the system has irregular agents, the system is still able to remove them completely.

4.7 Conclusions

HADES is a cancer inspired multi-agent system that is able to control and protect itself via basic protocols and a rescue communication protocol. Its basic protocols control the cloning, repair, movement, and self-induced death that govern each agent in the system. These protocols would be sufficient to control HADES if errors were not possible. However, since each cloning event has a probability of mutation, the system must be fault tolerant to deal with completely mutated agents. The rescue protocol has therefore been designed to allow agents to influence the death of neighbors due to violation of spatial buffers, enabling the system to self-maintain despite irregular agents. It allows agents to send anonymous messages within a radius to all other agents, mimicking chemical diffusion seen in nature. Our mechanism of inducing self-death/removal is designed such that it could be applied to any self-regenerating system, making it a valuable tool for multi-agent systems.

During the design of the rescue protocol and numerous simulations we found that in order to fully extinguish the irregular agents, the protocol should include two kinds of signaling: PLEASE DIE and I'M DYING . In the first an agent that recognizes irregularity surrounding it sends PLEASE DIE messages, and in the second an agent that is going to die announces its death to the environment as a way of transferring the alert for irregularity to its neighbors. Without this combined set of signals, the system would not be able to remove an entire cluster of irregular agents while maintaining the overall system health. We tested the system on a variety of basic protocol parameters, as well

as rescue protocol parameters. The results show that even as we give the system more difficult situations where irregular agents have increasing advantages over normal agents, the system is still able to remove all of the irregular agents at least 60% of the time. In the easiest case where irregular agents only clone themselves six times more frequently than normal agents, the system is able to remove all irregular agents 100% of the time in most cases. The results also show that the rescue protocol can remove all irregular agents consistently, given the correct rescue protocol parameters (set five).

HADES was also perturbed by allowing irregular agents to fail by either ignoring received messages or not sending the I'M DYING message upon removal from the system. In both cases the system is able to remove the irregular agents as long as they follow the protocols approximately 30% of the time. The system is also able to remove all irregular agents even when the rescue protocol is not active until 50% of the system is filled with irregular agents. Thus the rescue protocols are robust to failures within itself, and we do not need to assume that irregular agents always follow those protocols either.

Although our system is inspired by biology, the solution is designed for general multi-agent systems with citizenship where the agents share the goal of keeping the system functioning. As an autonomous sensor network, [16] is an example of this type of system. The improved sensor networks will be able to determine if incorrect data is initiated from a specific group of sensors, and thus send messages to shut them down. These sensors can then be repaired and re-introduced into the system. Another application for our rescue protocol is improving distributed software by adding real time repair. For instance, the system may have a process that runs concurrently to its other processes that enables it to declare fellow processes as damaged. This repair can be in the form of killing the process and then restarting a new one, as in HADES. A system that corrects C and C++ code as it runs was proposed in [102]. It does so by replicating the programs and comparing results from each replication, only using results with a majority agreement.

Repair mechanisms can also facilitate self-organization in agent systems as demonstrated in [120, 119]. A group of cooperative autonomous agents working toward a system level goal can benefit from repair controlled self-organization. For instance, if a single agent's decision making fails it will put pressure on the system by providing incorrect output that could damage the system's organization. This output can eventually be repelled by the other agents via the rescue protocol,

enabling the agents to re-organize the system after resetting or re-creating the faulty agent. We can improve on many systems that require continuous functioning in this way, as long as the control structures have relative autonomy. We therefore use simple rules of communication to respond to faults within a multi-agent system such that the system is able to recover from those faults.

CHAPTER 5

EMOTIONS FOR PREDATOR PREY

5.1 Introduction

Collective behavior can refer to both human and animal tendencies to influence each other's behavior. In this chapter we are interested in collective animal behavior that leads to improved species survival within a predator-prey system. A major topic of population dynamics, the study of the development of either a single or multiple interacting species, is the cycling of predator and prey populations. Predator-prey dynamics relate to a wide variety of ecological situations, from microbial phagocytosis to lions and gazelles. Most often predator-prey systems are built to describe animal species, with at least one species as prey and one as predator; however, they are not limited to describing only two species. The Lotka-Volterra [86] equations are commonly used to model this type of interaction, and are based on the classic logistic equation. However, it has been argued that these equations are not sufficient for modeling natural phenomena [83].

Cellular Automata (CA) offer a popular mechanism to analyze population dynamics as they directly represent spatial interactions between entities [71]. CA allow the creation of rules for determining how an entity will interact with its neighbors. The most popular version of a self-regenerating cellular automaton is the Game of Life, developed by Conway [55]. In the Game of Life cells are created or removed for the next time step based on the number of neighbors the cell has in the current time step. Although the rules can be completely defined in a single sentence, the dynamics are complex and still not completely understood. This ability of CA to give rise to complex dynamics via simple rules enhances its desirability for modeling complex phenomena, assuming that the appropriate simple rules can be designed. Thus, in population dynamics models, entities can explicitly exist on a grid and interact with specific neighbors. The system not only knows how many of each species is in the system, but to what extent they are mixed. The world can either be viewed as a torus with periodic boundary conditions or a bounded box that may or may not be square. A torus is beneficial for analysis and computation as all cells have the same number

of neighbors. However, in many ways a bounded region is more realistic, as the ecosystem of a set of species will not extend completely around the world but instead exist in some localized area.

We introduce intra-species disease transmission and emotion-inspired rules for our predator and prey (foxes and rabbits) model. Real populations in nature are subject to epidemic diseases, a number of which can cross species. Such diseases have significant effects at the level of individual behavior and population dynamics. Evidence suggests that a primary contributor to the evolution of the emotion disgust is protection from the risk of disease [34]. We explore the relationship between disease transmission and emotional response. For collective behavior to arise information is shared between conspecifics (members of the same species) and individual decisions are made on that information. The information shared is in the form of emotions, and both rabbits and foxes make decisions considering their emotions and the shared emotions. Emotions are affected by environmental events, and thus represent a high level of information about the environment. The development of emotions in higher animals has been conjectured to originate for purposes of survival in basic scenarios such as predator-prey [15, 87], and thus emotionally-inspired rules are a natural extension to the traditional CA framework. Although they have been suggested previously for CA [3] and have the potential to increase our ability to accurately depict changes within an ecosystem, we are unaware of any work utilizing emotions in the context of predator-prey dynamics modeled within a CA framework.

Similarly, in the last few years it has been suggested that emotions constitute an important part of adaptive decision making systems, contradicting the older view that emotions typically interfere with decision making [36, 132]. Case studies reported that people who suffered injury to or loss of areas of the brain related to emotion also experienced impaired decision making [13]. Instead of showing the benefit of decision making to only an individual when emotions are involved as is done in these studies, we examine the benefit of emotions for the group.

Thus, we choose to include the six basic universal emotions defined by Ekman [45] to our rabbits and foxes: happiness, sadness, fear, anger, disgust, and surprise. Emotions occur in response to specific world events, such as the happiness of food consumption and the fear of predator encroachment. Additionally, we enable conspecific communication of emotions to aid in coordination and cooperation. In other words, the emotional state of a member of a species will be communicated to a member of the same species within a restricted surrounding, and affect their emotional state.

We consider this approach to emotional communication as an efficient way of transferring information that is crucial for the survival of the group, and analyze two different techniques for sharing emotion: direct sharing from neighbors, and stigmergic sharing [98]. Our goal is to determine which modeling approach leads to the best collaboration within a species. Although our individuals are inherently selfish in that they make the best decision for themselves, our results show that communication of emotion can increase collective behavior for both predator and prey. Additionally, we show that it is in the best interest of both rabbits and foxes to use emotion when they do not know if the other species will use emotion, and explain this conclusion in terms of a population dynamics version of the classic Prisoner's Dilemma problem.

In this chapter we first describe related work in the fields of population dynamics, cellular automata, and emotions. We then describe our model in detail, including the two forms of communication and our evaluation mechanisms. Finally we discuss results and conclude.

5.2 Previous Work

Cellular automata have been modified to simulate and analyze many topics, with one of the first being von Neumann's description of self-replicating automata [152], which has been built upon extensively. One such extension is the work of Petraglio toward creating a cellular automata capable of performing arithmetic operations by using self-replication of the cells [115]. They have also been utilized for biological modeling of ants, where the floor is a set of CA controlling the ant movement "above" them [125].

Although these extensions maintained the traditional uniformity of all cells conforming to the same rules, CAs have also been extended such that each cell can run a different set of rules [140]. In Sipper's work this extension is used to evolve cellular automata capable of performing their desired computations. Additionally, all cells do not necessarily have to be updated synchronously, i.e. with all cells updating in parallel. However, the results of the system will vary based on whether or not updates are synchronous [134]. Synchronous and asynchronous updates can also be combined, where many different asynchronous updates are performed within a single synchronous step [62]. Arguments exist for the structure imposed by both synchronous and asynchronous CAs [73, 134]. Step-driven methods for asynchronous updates include different strategies for sweeping the cells to perform updates. Directional sweeps provide a simple strategy whereby cells are updated according

to their spatial order. Alternative sweeping techniques include fixed random sweeps and uniform choice, both randomization strategies. Both strategies are similar with the primary difference being replacement, which does not occur in the former, but does in the latter. A discussion of the practical implications of update ordering is presented in the next section.

An additional extension to CA is for population dynamics modeling, as CA allow the spatial environment to be directly modeled and thus for each entity to have specific neighbors with which to interact. The interaction of species living in symbiosis can be modeled with CA [4], as well as the growth and death of single species [148]. Typical results exhibit dynamics of populations growing and declining over time.

Predator-prey models often depend on spatial interaction and thus can benefit from CA modeling. Multiple interacting species move around the grid, with predator chasing prey. In some models, movement is purely defined by the birth of new entities into neighboring cells [39, 48], whereas other models allow individuals to actively move around the grid [40, 67]. We will allow two of our species to actively move (rabbits and foxes), and one to only move by reproduction (carrots).

Predator-prey models are usually based on the Lotka-Volterra equations [86], and thus they will allow for reproduction and death. Often the prey's need for food is either ignored or empty space is considered food [40, 48, 39]. It is also possible to require that reproduction occur during the same time step as eating [67]. Although a classic CA model would be completely based on neighbors for deciding actions, most predator-prey models are probabilistic [40]. The interactions between predator and prey can be greatly influenced by how prey chooses between avoiding predators and finding food [14].

For predator-prey dynamics in a CA it can be useful to analyze the patterns created in the system. This can include how the number of predators and prey fluctuate in the system over time with differing parameters, as well as how mixed they are spatially [39, 67]. Results from [39] indicate that a combination of Lotka-Volterra, an individual's ability to change, and the spatial structure of the CA give rise to both predators and prey self-organizing into self-sustaining patterns. It has also been proposed that taking the environment into account significantly affects results, potentially making them more realistic as living creatures naturally have outside influences other than a predator or prey [48].

CA models do not always have a single entity in each cell, although that is the most common approach [40, 67]. It is also possible to superimpose several layers, each corresponding to a different entity, similar to a population dynamics meta-population model [42, 48]. In this case there will be fluctuations in the percentage of a cell that is each type of entity. It is unclear whether this approach is more or less realistic than a single entity per cell model, although this multi-layer structure allows an easier approximation of differential equations to model fluctuations within each cell [42, 48]. However, it has been argued that a CA model can provide better results than a partial differential equation model [67], and they are viewed as the norm for spatial predator-prey models.

A goal within the predator-prey modeling community is to move toward more realistic models. One possible previously unexamined direction is to include emotion-based features for each species in the environment. Although there are a number of human psychological theories of emotion [45, 118, 124], it is generally agreed that emotions serve the purpose of increasing our ability to interact with our environment in a successful manner. CA have been studied with emotions in the past for investigating the behavior of an individual. For instance, the interaction of individual emotions within a single person was studied using a CA to determine how they work together to influence an individual's behavior [3]. Emotions have also been studied as part of an artificial entity modeled with a CA [38]. The goal of our study is to introduce emotions to a species by including rules on how each individual gains and updates emotions, how the emotions of the individuals are shared with other conspecifics, and how each member changes its behavior locally based on its emotional state. The outcome is a coordinated group of individuals that act like a complex system with shared emotions. Such experimentation may be of importance both in understanding evolution of competing species as well as in coordinating multi-robot systems.

Rolls has argued that human emotions have ten functions, including reflexive behaviors and motivation [124]. These functions are advantageous for predators and prey as well, causing them to react quickly when near the other. It has been suggested that human emotions were initially evolved due to the need to survive, showing many commonalities with both the reaction of prey to predator and predator to prey [118, 87]. Although there have been arguments that any defensive action in a prey represents emotions, more recently that view has been modulated to instead argue that although this may not always be the case, it is still likely that initial reactions to threat in animals can precede emotions [15].

Collective animal behavior is often studied from a few different perspectives: examining biological causes of collective behavior in animals [98], examining how to build models to capture collective behavior [31], and being inspired by collective behavior [18]. Often these overlap, as more recent interdisciplinary efforts use modeling to help decipher the biological phenomena and create a better understanding of what underlying mechanisms can result in various forms of collective behavior.

We therefore study computational emotions in the predator-prey context via a CA by including rules on how each individual gains and updates emotions, how the emotions of the individuals are shared with other conspecifics, and how each member changes its behavior locally based on its emotional state. Groups of prey in real situations will exchange information about their surroundings, increasing their likelihood of survival [141]. In the proposed model, the communicated emotions will include hints about a variety of survival conditions such as satisfaction (from food), fear (from predator), and disgust (from food poisoning). These emotions will result in behavior directing the individual to move in the best direction for survival. We hypothesize that the outcome will be a coordinated group of individuals.

5.3 Our Model

We use a four-species CA model to examine predator-prey dynamics in an environment that represents disease and where emotions are developed individually and communicated to neighboring conspecifics. Each point in the model may hold a single entity from any species at any given point in time. For ease of description we will label our primary species as rabbits, foxes, and carrots. All are able to reproduce and die, and foxes and rabbits are also able to move. Empty spaces (the fourth species, vacancies, as described in the literature) represent an area where the other species may reproduce or move. Our grid structure is defined as a torus, and thus there are no corners requiring special treatment. Each individual's next step is determined based on probabilities and their neighbors within the Moore neighborhood (8 neighbors), as suggested to be ideal by [27]. The implementation of the CA is asynchronous. The model world is shown in Figure 5.1 to demonstrate the non-uniformity of the distribution of entities within the simulation. By using stochastic movements on a toroidal grid the potential for structure induced by directional sweeps is essentially removed.

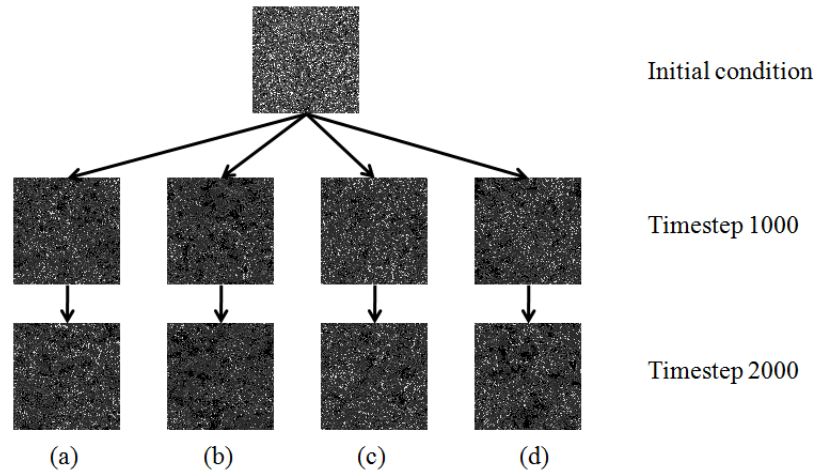


Figure 5.1. Simulation world at multiple time steps shows the dynamics of model, and how it may differ between a world (a) without emotion, (b) without fox emotion but with rabbit emotion, (c) with fox emotion but without rabbit emotion, and (d) a world with fox and rabbit emotion. Vacant squares are in black, carrots are in dark gray, rabbits are in gray, and foxes are in white.

At the start of the first simulation, foxes, rabbits, and carrots are placed randomly such that no individuals overlap. Carrots are food for rabbits, rabbits are food for foxes, and foxes are at the top of the food chain. For simplicity all reproduction is asexual. Some aspects of the model results may be slightly less realistic as a result, but as the goal is to examine the use of a modeling technique instead of directly modeling a specific environment, our hypotheses can be equally tested in either case. A predator eats a prey when it moves to the prey's position. Once all of a species have been eaten, no more of that species can come into being in the system (Figure 5.2(a)). However, as long as there is still at least one of a species alive it is possible for it to reproduce to create a new one (Figure 5.2(b)).

Both rabbits and foxes may starve, and therefore the number of carrots in the model can strongly affect the overall dynamics of the system. We choose an adequate reproduction rate and initial number of carrots such that rabbits are unlikely to starve due to inability to find food. Each entity is also capable of becoming diseased. Disease initiates in the carrot population, and moves to the rabbits and then the foxes when eaten. A diseased rabbit or fox will become hungry at double the rate of a healthy rabbit or fox, thus increasing their chance for reaching starvation and dying.

Rabbits and foxes can exist either with or without emotions. Carrots are not affected by emotions. We will first describe the model without emotions, and then describe how emotions are gen-

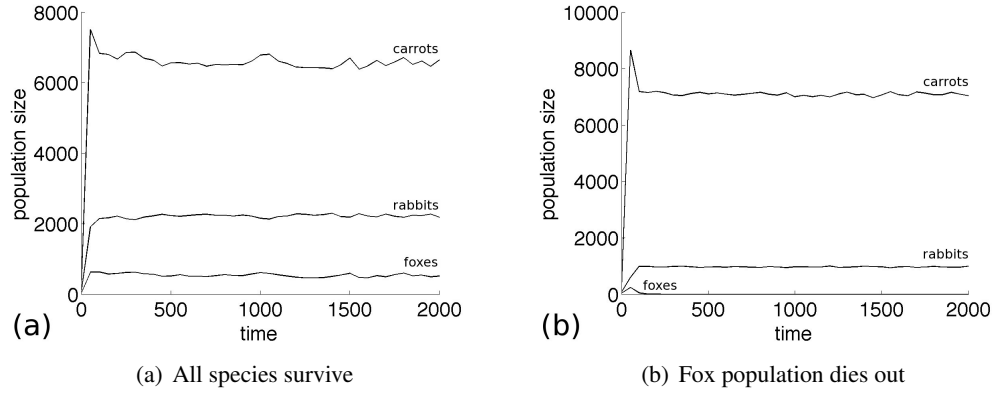


Figure 5.2. Population dynamics of the system with different parameter settings. In (a) all population sizes tend towards non-zero attractors. In (b) the fox population crashes while the rabbit and carrot population sizes continue to tend towards non-zero attractors.

erated and how they modify individual behaviors within the model. In our results we will compare the scenario of no emotion with the scenario of using emotion.

5.3.1 Probabilistic and Neighbor-based Rules

At each time step, each cell occupied by a carrot will update following these rules:

- **Reproduction:** If it is at maturity age it produces a new carrot into a vacant adjacent square, if one exists. Age is reset to zero.
- **Disease:** When a new carrot is produced it has a 0.1 probability of being diseased. If the parent is diseased, the probability doubles.
- **End of disease:** Disease lasts for a minimum of 2 time steps. After that minimum time has elapsed, there is a fixed probability $Cure$ of disease being cured each time step.

Rabbits and foxes follow slightly more advanced rules at each time step (Figure 5.3):

- **Movement:** Movement occurs by computing a local gradient of preference for the surrounding cells as seen in Equation 5.1. Species can only move to a cell containing another individual if that individual is their food source.

$$g_{d,sp}(t, x, y) = food_{sp}(t, x_d, y_d) \quad (5.1)$$

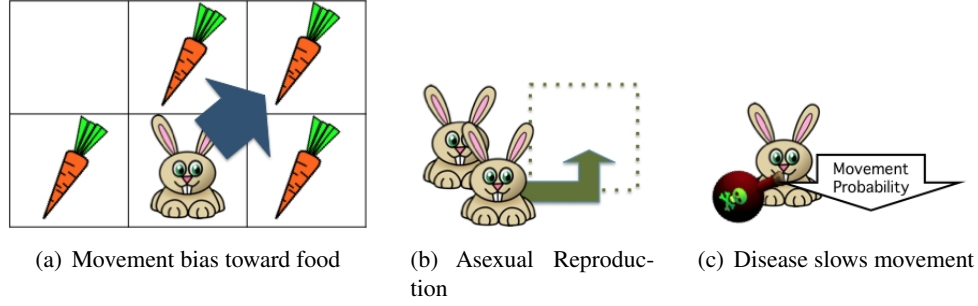


Figure 5.3. Probabilistic and Neighbor-based Rules for rabbits and foxes.

where d in $D = \{NW, N, NE, W, E, SW, S, SE\}$, t is the current time, $g_{d,sp}(t, x, y)$ represents the preference for each direction d at time t for an individual at location (x, y) , $sp \in \{\text{rabbits, foxes}\}$, and $food_{sp}(t, x_d, y_d)$ returns the existence of prey for a given species, sp , at locations in direction d from position (x, y) at time t calculated on that grid point and its two neighbors. These preferences are converted to normalized probabilities that bias the individual's otherwise random movements.

- **Movement rate:** When diseased it will have a decreased probability of movement each time step (*DiseaseMove*).
- **Hunger:** Hunger increases by 1 each time step it does not eat, or by 2 if it is diseased. Hunger is decreased to zero when it eats.
- **Disease:** If a predator eats a diseased prey, the predator becomes diseased as well. Disease lasts for a minimum of 2 time steps. After that time elapses there is a probability *Cure* of disease being cured each time step.
- **Reproduction:** If it is of maturity age there is a fixed probability *Rep* of a new rabbit/fox being created in a vacant adjacent square, if one exists. The individual's age will continue to increase until it reproduces, and then it will be reset to zero.

After all entities have followed these steps, each will increase their age by 1.

5.3.2 Individual Emotions

Rabbits and foxes can use emotion in the model, but carrots do not. Emotions are calculated at the end of the sequence described in the previous subsection, and are used by rabbits and foxes

Emotion	$X_{e,sp}(t,x,y)$
Happiness	1 if ate prey, 0 otherwise
Sadness	$t - [\text{timestep of last reproduction}]$
Anger	exp^{hunger}
Fear (fox)	Anger of neighboring foxes
Fear (rabbit)	Number of neighboring cells with foxes
Disgust	1 if ate diseased prey, 0 otherwise
Surprise	$\sum_e \frac{E_e(t,x,y) - E_e(t-1,x,y)}{5}$, where e does not include surprise

Table 5.1. Experience values affecting each emotion. For each emotion the corresponding value of the experience variable is listed.

when determining the next movement direction as well as their probability of reproducing. They are based on Ekman’s original six basic emotions (fear, anger, sadness, happiness, disgust, surprise; [45]). Each individual maintains values of their own emotions. Emotions are independent for both rabbits and foxes, and each emotion is affected by different experiences related to neighbors, hunger, reproduction, and disease. For rabbits, fear increases with the number of surrounding foxes. For foxes, fear increases with the amount of anger felt by surrounding foxes. Anger for each species increases exponentially based on hunger level. Sadness increases the longer the individual has gone without reproducing. Happiness increases after food consumption. Disgust increases when an individual is diseased. Surprise increases by the average amount of change in all of the other five emotions from one time step to the next. The numerical values used to represent these experiences when calculating emotions can be seen in Table 5.1. Individuals also communicate their emotion, which can influence the emotions of other nearby individuals from their own species. Communicated emotions of one species cannot be seen or interpreted by the other species.

An individual’s emotions are based on experiences, as well as their previous emotion and the emotions being communicated nearby. They are computed at each time step as seen in Equation 5.2. The previous emotion and communicated emotion are both discounted, to prevent them from overpowering newer experiences or causing monotonically increasing emotions over time.

$$E_{e,sp}(t+1, x, y) = (1 - c_{m,sp})(X_{c,sp}(t, x, y) + c_{c,sp} * CE_{e,sp}(t, x, y)) + c_{m,sp} * E_{e,sp}(t, x, y) \quad (5.2)$$

where $e \in \{\text{fear, anger, sadness, happiness, disgust, surprise}\}$, $sp \in \{\text{rabbits, foxes}\}$, $X_{e,sp}(t, x, y)$ is the unique experience of each emotion for each species (Table 5.1), $c_{m,sp}$ is the memory dis-




Figure 5.4. The two communication styles. The 'x' represents where an entity currently exists, and the colored background represents the currently shared emotion at that cell (darkest is strongest). Both images represent the state of the shared emotion from that entity after three time steps.

counting coefficient that determines what percent of the new emotion is based on new versus old emotional information and is calculated as in Equation 5.3, $c_{c,sp}$ is the discounting coefficient for communicated emotion, and $CE_{e,sp}(t, x, y)$ is the communicated emotion at position (x,y) at time t as shown in Equation 5.4 and Equation 5.5. The coefficient $c_{m,sp}$ is bounded to $[0.1, 0.5]$ and $c_{c,sp}$ is bounded to $[0, 1]$.

$$c_{m,sp} = 0.1 + 0.4 * \frac{abs(E_{surprise,sp}(t, x, y) - E_{surprise,sp}(t-1, x, y))}{E_{surprise,sp}(t, x, y) - E_{surprise,sp}(t-1, x, y)} \quad (5.3)$$

Emotions are computed for each individual. The calculation of communicated emotion is dependent on which communication paradigm is in use.

5.3.3 Emotion Communication - Direct

In direct communication, only emotions from the immediate neighborhood from the previous time step are used in the decision (Figure 5.4(a)). Additionally, emotions are only read from the map if a conspecific existed at that point; i.e., only emotion shared from events in the previous time step are included. An individual will overwrite old emotions at their location between time steps. When an individual is determining their own emotions for time step (t+1), they will calculate what emotions are being shared through the environment via Equation 5.4 from time step t.

$$CE_{e,sp}(t, x, y) = \sum_{d \in D} S_{sp}(t, x_d, y_d) * E_{e,sp}(t, x_d, y_d) \quad (5.4)$$

where $e \in \{\text{fear, anger, sadness, happiness, disgust, surprise}\}$, $sp \in \{\text{rabbits, foxes}\}$, D represents all directions, $S_{sp}(t, x_d, y_d)$ returns 0 or 1 denoting the existence of species sp in direction d from position (x,y) at time t , and emotion $E_{e,sp}(t, x_d, y_d)$ is the amount of emotion in direction d from position (x,y) at time t . After emotions and communicated emotions have been calculated for all entities in the system, all emotions (both internal and communicated) are decayed linearly by a small value denoted in Table 5.2.

5.3.4 Emotion Communication - Stigmergic

With stigmergic communication, the shared emotions will remain at the cell over time and will decay with each time step (Figure 5.4(b)). When an entity moves to that cell and leaves its own emotional mark, the newly shared emotions will be combined with the previously shared emotion. This thus affects the calculation of what shared emotions an individual's emotion is affected by, as seen in Equation 5.5.

$$CE_{e,sp}(t, x, y) = \sum_{d \in D} E_{e,sp}(t, x_d, y_d) \quad (5.5)$$

where $e \in \{\text{fear, anger, sadness, happiness, disgust, surprise}\}$, $sp \in \{\text{rabbits, foxes}\}$, D represents all directions, and emotion $E_{e,sp}(t, x_d, y_d)$ is the amount of emotion in direction d from position (x,y) at time t .

After emotions and communicated emotions have been calculated for all entities in the system, all emotions are decayed. Although the individual emotion is still decays linearly as with the Direct Communication, we test two different form of decay for the communicated emotion: linear and geometric (values in Table 5.2). This variety of decay examines whether the rate of decay affects the emotional decisions of the predator or prey.

5.3.5 Rules Enhanced by Emotions

Rabbits and foxes with emotions have altered reproduction rates, and movement direction preferences from their unemotional counterparts (Figure 5.5). Other aspects are calculated the same as shown in the Probabilistic and Neighbor-based Rules section.




Figure 5.5. Some emotions modify fox and rabbit behavior negatively, and others modify the behavior positively. Both movement and reproduction are affected by emotion.

Reproduction Rates: Rabbit and fox reproduction rates are altered positively by happiness, and negatively by disgust and anger. Additionally, rabbits will not reproduce at all while their fear is above a threshold. The calculation of reproduction rate can be seen in Equation 5.6.

$$\begin{aligned}
 R(t) &= Rep * (1 - Rep) * (E_{happy}(t, x, y) \\
 &\quad - Rat * E_{disgust}(t, x, y) - (1 - Rat) * E_{anger}(t, x, y))
 \end{aligned} \tag{5.6}$$

where $R(t)$ is the probability of reproducing at time t , Rep is the initial probability of reproducing after the maturity age has been reached, Rat is the ratio of how much disgust versus anger decreases reproduction, and $E_{happy}(t, x, y)$, $E_{disgust}(t, x, y)$, and $E_{anger}(t, x, y)$ represent current emotional values.

Movement Preference: The local preference gradient for movement considers emotions, where an individual will move toward the highest positive value $g_{d,sp}(t, x, y)$. This differential is taken to be the difference between the emotion in a given direction and the current emotion of the individual, as seen in Equation 5.7.

$$\begin{aligned}
 g_{d,sp}(t, x, y) &= food_{sp}(t, x_d, y_d) \\
 &\quad + \sum_e valence_e * (E_{e,sp}(t, x_d, y_d) - E_{e,sp}(t, x, y))
 \end{aligned} \tag{5.7}$$

where (x, y) represents the individual's current location, (x_d, y_d) represents locations in direction d , e is taken over all emotions except surprise, and $valence_e$ is -1 for a negative emotion e (fear, anger, sadness, disgust) and 1 for a positive emotion e (happiness).

Emotions are used to encode and communicate various features of the environment to modulate the behavior of individuals. The components of the environment incorporated into emotions are intuitively useful for survival, which suggests that emotions should modify the behavior of individuals in a way that is beneficial for their species.

5.4 Results - Role of Emotion in Predator and Prey Decisions

We first examine how emotions shared via direct communication affect the dynamics between foxes and rabbits, specifically their ability to reproduce, avoid starvation, and avoid disease. In the next section we examine how varying the communication paradigm and amount of information shared with neighbors changes the population and emotions of both foxes and rabbits.

5.4.1 Experimental Design

The simulation is run on a grid world of size 100 x 100. Each point on the grid interacts with its Moore neighborhood of radius 1. Simulations are run for a total of 2000 time steps with an initial random placement of individuals on the grid. Each of the twenty initial placements are tested on the four emotion scenarios: no emotions, only foxes using emotions, only rabbits using emotions, and both species using emotions.

Initial population sizes for foxes, rabbits, and carrots were 1000, 2000, and 6000. Initial parameter searches were done on both population parameters and emotion parameters for all populations. The initial parameter searches for emotion were done for when only rabbits have emotion or only foxes have emotion, to determine how the individual parameters affect the overall population dynamics. From these parameter searches the most promising parameters were taken to combine and investigate further (Table 5.2 and Table 5.3).

We examine our claim that *computational emotions provide a framework for modeling predator-prey dynamics that will provide different modeling behavior than a traditional cellular automata model* by examining a series of related hypotheses:

Hypothesis 4.1 *The use of disgust will decrease a population's overall rate of being diseased.*

Hypothesis 4.2 *The communication of emotions will increase population size for a species.*

	Carrot	Rabbit	Fox
Rep	n/a	0.75	0.75
Maturity Age	2	6	7
Cure	0.8	0.8	0.8
DiseaseMove	n/a	0.8	0.8
Starvation	n/a	2	6

Table 5.2. Non-emotion parameters for each species. Rep is the probability of reproducing after reaching the Maturity Age, Cure is the probability of being cured from, DiseaseMove is the probability of an individual moving if it is diseased, and Starvation is the hunger level that causes death.

	Fox	Rabbit
Decay rate	0.45	0.45
<i>Rat</i>	0.4	0.4
$c_{c,sp}$	0.55	0.75
Fear threshold	n/a	0.5

Table 5.3. Emotion parameters used in experiments. Decay rate linearly decrements the emotion value at every point, Rat denotes the ratio of how disgust and anger affect reproduction rate, $c_{c,sp}$ discounts communicated emotion from surrounding conspecifics, and the fear threshold denotes the level of fear necessary to pause rabbit reproduction.

Hypothesis 4.3 *Fast decay of shared emotions will result in population sizes most closely resembling a population with no emotion.*

5.4.2 Population Dynamics

The basic dynamics of the system can be seen in Figure 5.6 for the four scenarios: without emotion, only foxes having emotion, only rabbits having emotions, and both populations having emotions. In Figure 5.6(a) we see that rabbits benefit the most if they are the only population using emotions, and suffer the most when foxes are the only population using emotions. Thus, rabbits have the most evolutionary benefit when foxes do not use emotions, whether or not rabbits use emotion. As will be explained by later figures, this disparity is due to emotions enabling foxes to eat rabbits much quicker, and emotions only marginally increasing a rabbit's ability to flee from foxes. This directly impacts the carrot population (Figure 5.6(c)): the ordering of best to worst scenario for carrots is a direct inverse of the population size ordering for rabbits.

Fox-only emotion gives the fox the best situation, and rabbit-only emotion gives them the worst situation (Figure 5.6(b)). The scenarios of both populations having emotion or neither having emo-

tions are in the middle, although with no significant distinction between them unlike for rabbits. Emotions increase a fox's ability to find rabbits to eat, but do not decrease its disease rates. The following figures will examine this disparity further.

Population changes are directly affected by death and reproduction. The causes of death are highly inter-related and thus should not be examined separately. For rabbits and foxes death is caused by starvation, which is less likely with higher food consumption, and more likely with higher disease rates. Rabbit death is also caused by foxes eating them, so fox consumption is related to rabbit death as well. Reproduction is highly affected by hunger, disease, and nearby predators due to emotions. A rabbit will pause reproduction when it has high fear, and both rabbits and foxes reproduce less frequently when they have not eaten recently, or have recently eaten diseased food. Thus, we must examine death and reproduction for rabbits and foxes to fully understand the population changes.

5.4.3 Rabbits

Rabbits eat most frequently (Figure 5.7(a)) and are the most diseased (Figure 5.7(b)) when foxes use emotion. Disease is highly correlated to the amount of carrots eaten and fox emotion. Rabbit-only emotion minimizes rabbit disease, so disgust is allowing rabbits to avoid diseased carrots. Rabbits may have the lowest population size when they eat the most due to less food being available for other rabbits. Rabbits starve at a very low rate overall, the fastest being with rabbit-only emotion, where they eat the least and are the most diseased (Figure 5.7(c)). The rabbit population is not highly dependent on rabbit starvation rates though, as the scenario of highest starvation (rabbit-only emotion) is also the scenario of highest population. Starvation rates are inversely related to population, and thus reproduction must play a vital role.

Rabbits reproduce the least when they use emotion (Figure 5.8(a)), likely due to fear, disgust, and anger significantly decreasing their probability for reproduction. In the highest population scenario of rabbit-only emotion, rabbits are reproducing the least, eating the least, having the lowest disease, and starving the most. Rabbits are probably best off with rabbit-only emotion because it allows them to escape foxes the easiest, as reproduction rates and death by starvation do not explain the population changes. This hypothesis is supported by Figure 5.9(a), where the rabbit-only emotion and no emotion scenarios are seen to decrease fox consumption.

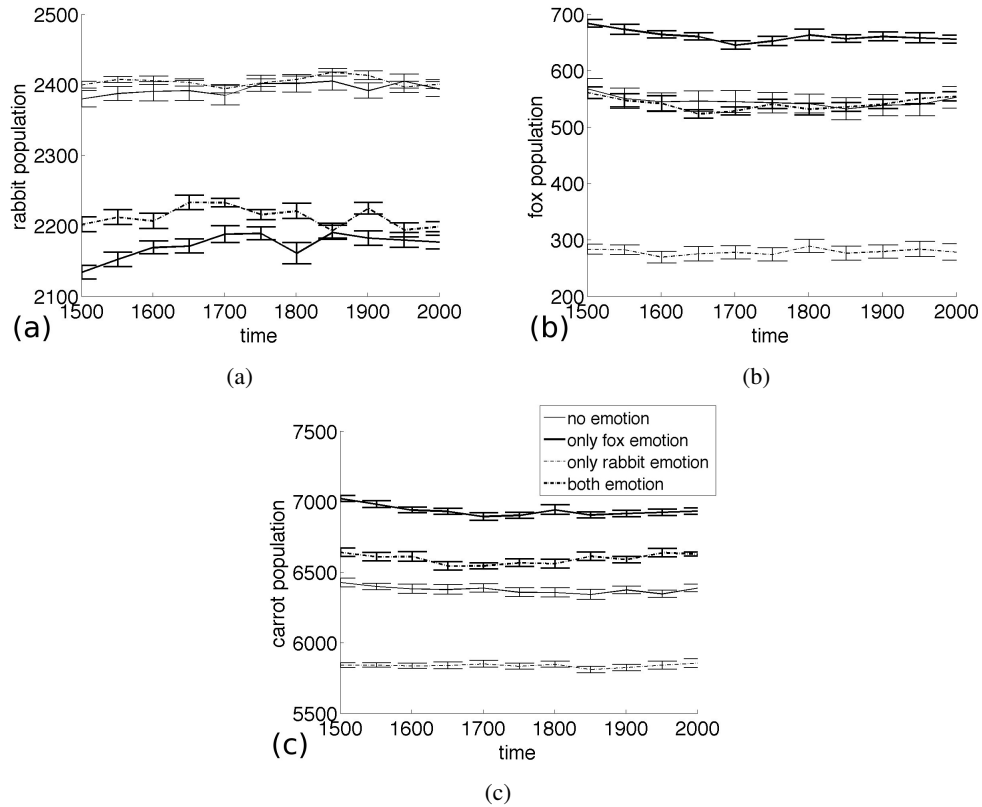


Figure 5.6. Population variation over time. The population over time is shown for no emotion, rabbit-only emotion, fox-only emotion, and both species emotion. The rabbit population (a) is highest when only rabbits have emotion with no emotion as a close second, and lowest when only foxes have emotion, and all emotion as second worst. The fox population (b) and carrot population (c) are both highest when only foxes use emotion, and lowest when only rabbits use emotion, with all emotion and no emotion close in the middle.

5.4.4 Foxes

The fox population is benefited the most by fox-only emotion, but no emotion and both rabbit and fox emotion scenarios are tied for second. Foxes consume the most rabbits when they use emotion to enable them to track rabbits (Figure 5.9(a)), but they are also the most diseased when they use emotion (Figure 5.9(b)). Fox disgust does not appear to be successful in allowing foxes to avoid eating diseased rabbits, although it is successful in allowing rabbits to avoid diseased carrots. Most likely this is due to a fox's food being mobile. If a fox is seeing high disgust to the south from other foxes, that does not mean that the diseased rabbits are currently in that direction, only that they recently were in that direction. Rabbits do not have this problem as carrots are stationary, so their diseased relatives will stay near where a diseased carrot was recently eaten.

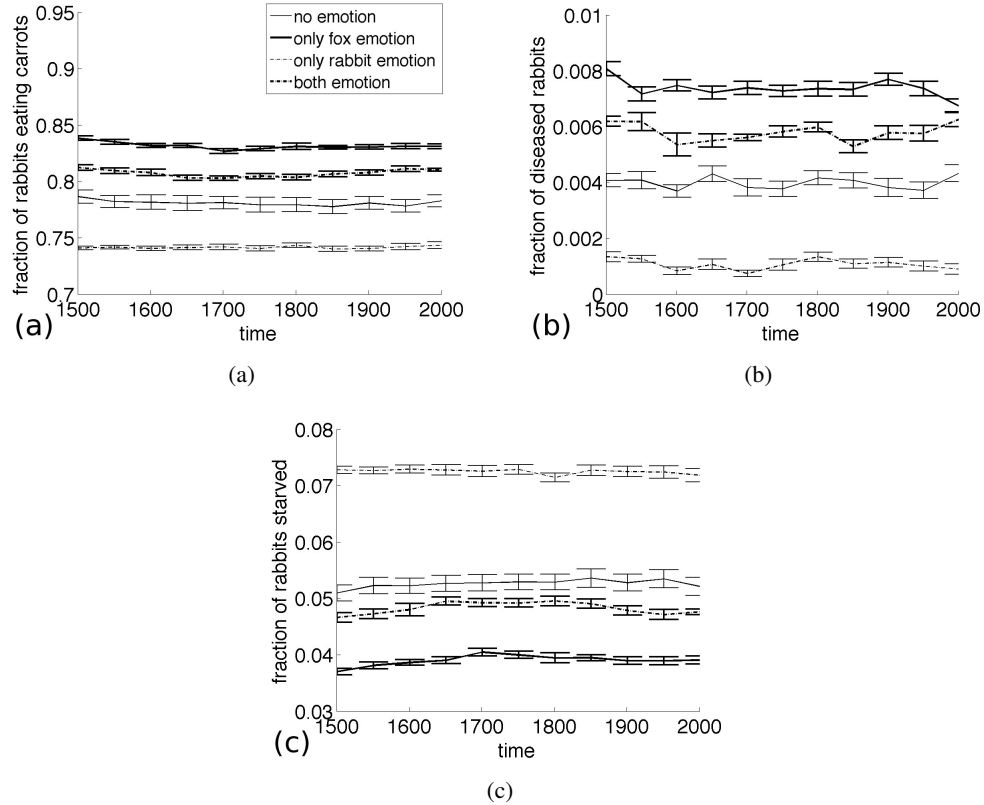


Figure 5.7. Rabbit consumption, disease, and starvation are all correlated. (a) Rabbits eat most frequently when foxes use emotion. This graph is the inverse of Figure 3(a), as rabbits eat the most when they have the lowest population size. (b) Rabbits are the most diseased with fox-only emotion, and the least diseased with rabbit-only emotion. Disease is highly correlated to amount of carrots eaten. (c) Although rabbits starve at a very low rate overall, they starve the fastest in the scenario where they eat the least and are the most diseased. These trends are very similar to the trends in Figure 5.6(a).

Foxes starve slowest when they use emotion, corresponding to when they eat the most and are thus the most diseased. Fox starvation is close to completely describing the trends seen in Figure 5.6(b), however the case of no emotion results in a higher population count despite also resulting in higher starvation. Therefore reproduction must play a significant role in population numbers as well. Foxes reproduce the most when they do not use emotion (Figure 5.8(b)). Anger and disgust are thus noticeably decreasing fox reproduction. Fox reproduction is not as significantly decreased by emotion as rabbit reproduction is since fox reproduction is not paused by fear. If we take both starvation and reproduction into account, however, the trends in Figure 3(b) are logical. Fox reproduction has a stronger affect on the fox population than rabbit reproduction has on rabbit

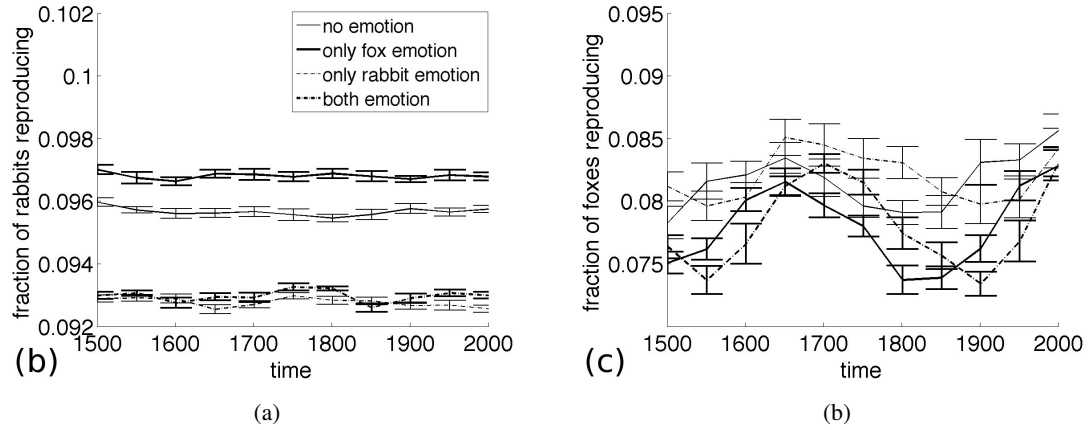


Figure 5.8. (a) Rabbits reproduce most frequently with fox-only emotion, closely followed by no emotion. Rabbits reproduce the least when they use emotion, likely due to fear, disgust, and anger significantly decreasing their probability for reproduction. (b) Foxes reproduce most frequently when they do not use emotion, most likely due to disgust and anger decreasing their reproduction probability. The effect is less pronounced than for rabbits as there is no fear level that stops them from reproducing completely.

population because there is no predator to counteract reproduction, so it is only slowed by lack of prey. This result is expected in predator-prey dynamics.

5.4.5 Discussion

Our introduction of emotion to a predator-prey model has shown multiple biases in the population dynamics: increased food consumption; reduced predation; and increased population sizes. These biases are not guaranteed by the introduction of emotion as the dynamics of the system links system attributes, i.e. too many rabbits can lead to overcrowding, decreased growth, and increased starvation. Nevertheless, the introduction of emotion to a species generally increases its population size. Additionally, since we see a loss in population size to a species when it does not use emotion but the other species does use it, it would be generally advantageous for a species to evolve emotions when it is unclear whether or not the other species is also evolving emotions. This is true despite the fact that both species using emotion causes an overall decrease in both population sizes.

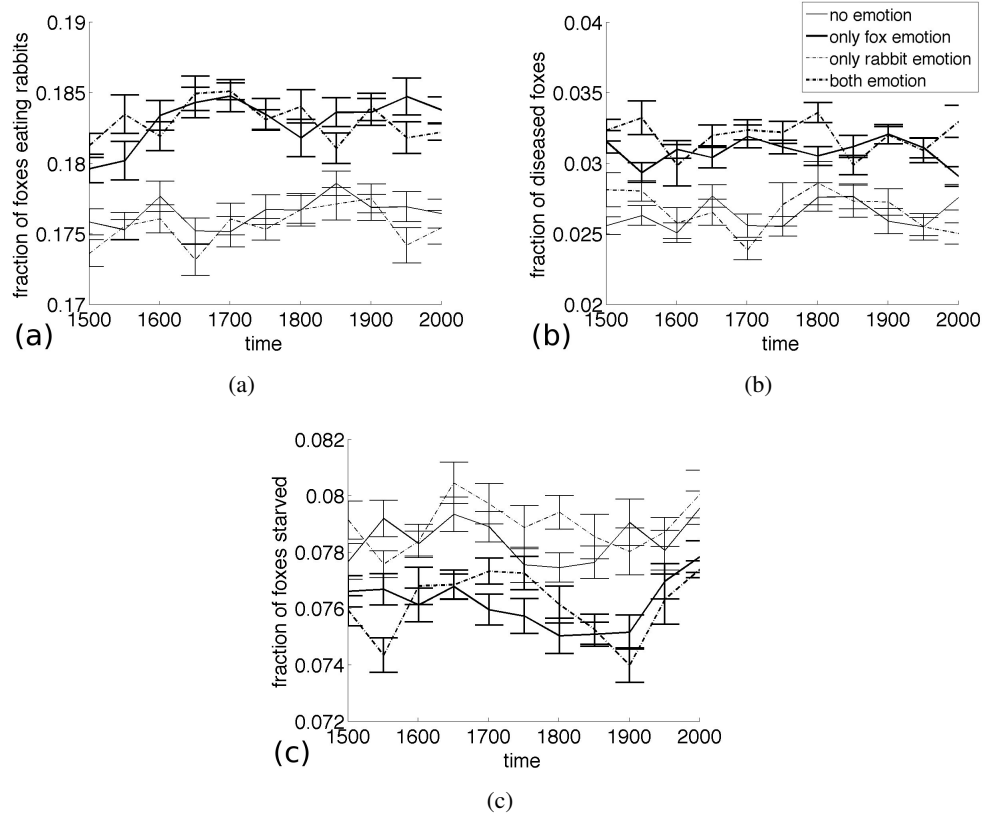


Figure 5.9. Fox consumption, disease, and starvation are correlated. (a) Foxes eat the most frequently when they use emotion. (b) Foxes are the most diseased when they use emotion, and thus when they eat the most. Thus, emotions are not improving a fox's ability to avoid diseased food. (c) Foxes starve fastest when they do not use emotion, corresponding to when they eat the least and are the least diseased. Starvation rate does not directly correspond to population size.

5.5 Results - Comparison of Communication Paradigms

5.5.1 Experimental Design

The simulation is run on a grid world of size 50 x 50. Each point on the grid interacts with its Moore neighborhood of radius 1. Simulations are run for a total of 2000 time steps with an initial random placement of individuals on the grid. Each of the twenty initial placements are tested on the four emotion scenarios: no emotions, only foxes using emotions, only rabbits using emotions, and both species using emotions.

Initial population sizes for foxes, rabbits, and carrots were 240, 500, and 1500. All parameters settings are the same as in the previous analysis, except for the communication coefficient $c_{c,sp}$ (Table 5.2, 5.3). We test three communication coefficients ($c_{c,sp}$: 0.1, 0.5, 0.9) for each communication

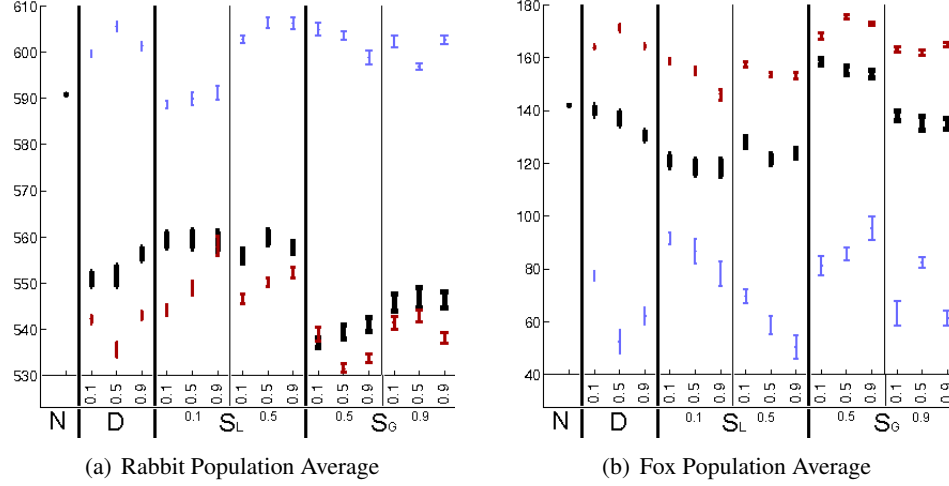


Figure 5.10. Average rabbit (a) and fox (b) populations for each parameter combination. The x-axis shows the communication coefficients (0.1,0.5,0.9) and decay values for each type of communication (N: no communication; D: direct; S: stigmergic (L=linear, G=geometric)). Black (darkest) represents when both species use emotion, Red (second darkest) represents when only foxes use emotion, and Blue (lightest) represents when only rabbits use emotion.

paradigm. For the stigmergic communication we test two types of decay, each with two emotion decay rates: linear (0.1,0.5) and geometric (0.5,0.9). We compare how the populations and individual emotions differ in terms of different communication paradigms (direct vs stigmergic), the amount of information used from the environment versus one's own information (communication coefficient), and how long the system and individual's memory of an emotion remains (decay rate).

5.5.2 Effect on Emotion

All rabbit emotions except for fear increase with an increase in the communication coefficient (Figure 5.11). Thus, there is enough emotion sharing (except for fear) on the grid that it increases the emotion when the environment has a strong effect on an individual. However, only a fox's fear is significantly increased as the communication increases; all other emotions are essentially constant (Figure 5.12). This suggests that there are too few foxes for effective communication, which is supported by the constant level of emotion when foxes use direct communication.

All rabbit emotions except anger and sadness are stronger when both species use emotion instead of only rabbits using emotion. The increase of fear in rabbits when foxes also use emotion is of particular interest, as it supports the idea that fox emotion improves their performance. Rabbit anger

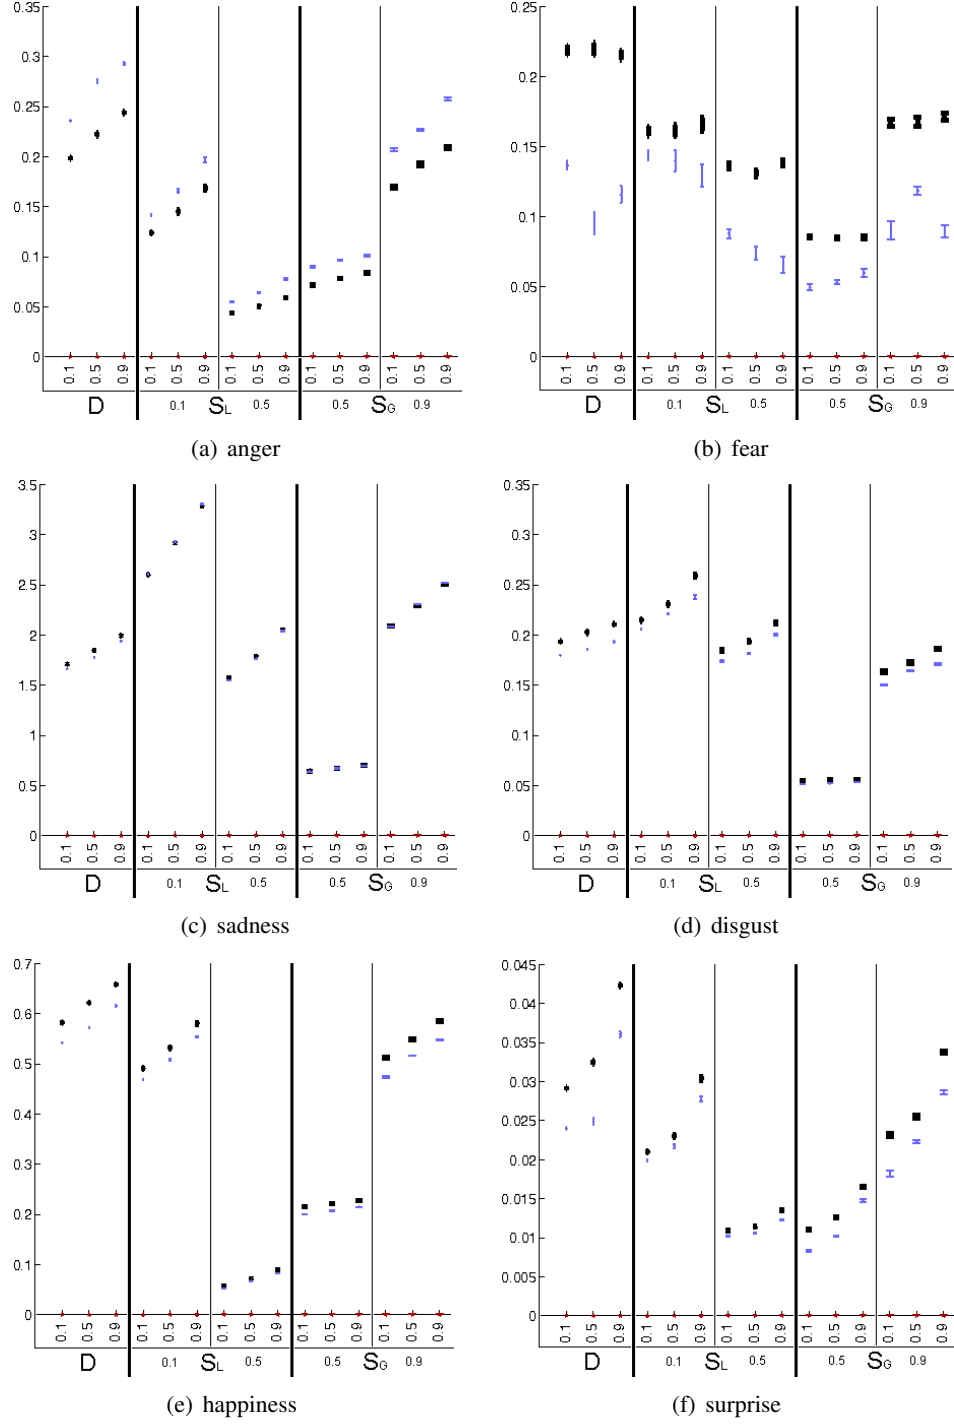


Figure 5.11. Average individual rabbit emotion for each parameter combination. The x-axis shows the communication coefficients (0.1,0.5,0.9) and decay values for each type of communication (N: no communication; D: direct; S: stigmergic (L=linear, G=geometric)). Black (darkest) represents when both species use emotion, Red (second darkest) represents when only foxes use emotion, and Blue (lightest) represents when only rabbits use emotion.

is increased when only rabbits use emotion, but rabbit sadness does not change. All independent fox emotions except fear are the same when either only foxes use emotion or both species use emotion.

For both species a higher geometric decay rate with stigmergic communication leads to more emotion than a lower geometric decay rate. A linear decay affects rabbit and fox emotion differently, however. For rabbits, stigmergic communication with a lower linear decay always leads to higher emotion than with a higher linear decay. For foxes, both linear decay rates lead to the same level of anger, disgust, and happiness. However, an increased linear decay rate leads to decreased sadness and increased fear for foxes.

Stigmergic communication with low linear decay leads to the highest levels of rabbit sadness and disgust. Rabbit fear is increased the most with either direct communication when both species use emotion, or with stigmergic low linear decay when only foxes use emotion. Linearly decayed stigmergic communication tends to lead toward higher fox anger, sadness, and fear than with a geometric decay; the opposite is true for disgust and happiness.

Direct communication leads to the highest levels of rabbit anger and happiness. Direct communication leads to lower fox anger and fear but higher sadness, disgust, and happiness than a linearly decayed stigmergic communication. Since anger and happiness are both related to hunger in inverse ways, it is not surprising that the fox emotion trends between the two emotions are essentially opposing each other.

All fox surprise trends essentially mimic fear. This leads to the suggestion that the primary factors of surprise, a composite emotion in our model, are fear and anger. As the emotional memory is inversely related to the amount of surprise experienced, fearful and angry foxes tend to base their memory and decisions upon more recent events. This is clearly a beneficial behavior that suggests that foxes will adapt their behavior for the environment when experiencing these negative emotions.

5.5.3 Effect on Population

Population averages are consistently better for each species when they are the only species using emotion (Figure 5.10). Foxes additionally benefit from emotion when both species use emotion and either direct communication is used (with less than 90% of communication coming from the environment) or geometric decay is used with stigmergic communication.

We expect that high decay rates in stigmergic communication would lead to results that are almost indistinguishable from the direct communication case, as emotions would not be able to linger long enough to be significantly different from only counting the previous time step. This is generally true with respect to the range in which the populations fluctuate while varying the communication coefficient for both linear and geometric decay in both populations.

The rate of decay for stigmergic communication generally has no effect on the rabbit population, only causing a significant difference in population when only rabbits use emotion with a high linear decay. The fox population decreases as decay rate increases when only rabbits use emotion, and essentially does not change when only foxes use emotion. When both species use emotion, however, fox population stays relatively constant as linear decay increases but decreases as geometric decay increases. A potential cause of this can be seen in the surprise and fear of rabbits. The increase in surprise at a high geometric decay will cause rabbits to focus more on the present than the past, while a higher level of fear will cause rabbits to flee foxes.

As communication coefficients increase most rabbit emotions also increase. This is to be expected as the increased coefficients should lead to more emotion present in the system. For this reason the relatively constant level of emotion for fear in rabbits when both species have emotion regardless of the communication coefficient is interesting. This suggests that there is an optimal amount of information for a prey to convey about predators.

Linear decay in rabbits may have an upper limit on the amount of decay, which is not surprising. Geometric decay may also have an upper limit for rabbits, with respect to the communication coefficients. This could be the point at which there is effectively no more communication in the system, or perhaps just for some of the emotions.

Geometric decay has more complex effects on the fox population, while linear decay seems to generally be detrimental. More communication without geometric decay is generally bad for foxes.

5.5.4 Discussion

We have analyzed the use of computational emotions toward increasing collaboration and collective behavior for both predators and prey in a Cellular Automata predator-prey model. Both species were given emotions inspired by Ekman's six basic universal emotions and current research

on how emotions affect animals in predator-prey scenarios. Two methods of communication were tested: direct communication, and stigmergic communication.

From the results we see that both species are benefited the most if they are the only species using emotion. However, we also see that the fox population benefits when both populations use emotion if information is stigmergic communication is used with a low geometric decay. Thus, the predator is able to act collectively with a trail of information when it remains for a longer period of time.

All rabbit emotions except for fear are increased by an increase in the communication coefficient. However, only a fox's fear is significantly increased as the communication increases; all other emotions are essentially constant. Thus, rabbit emotions are more strongly affected by older information, whereas fox emotion is not.

5.6 Conclusions

Predator-prey dynamics are frequently modeled by cellular automata due to the spatial ordering of entities within the system. This spatial ordering allows each member of the population to exist in a specific place on the grid and interact with its neighbors, potentially giving a more realistic dynamic among individuals. As some emotions have been found to evolve for survival in a predator-prey environment, we enhance the model by adding computational emotions based on Ekman's six basic emotions to our predator and prey. Conspecific communication of emotion allows individuals to transmit relevant local information to other members of its species.

Representing a biological system as a model always leaves some features of the environment unexplained or oversimplified. This has been true ever since the first use of physical laws to describe the real world. The use of cellular automata to model a biological system is no different. CA approximate decisions based upon complex state information with simplified rule sets. In studying this model it is important to consider that the model describes approximations of behaviors for rabbits, foxes, and carrots. It is also important to consider that the emotions are the authors' interpretations of previous studies and observations. Nevertheless, the model is constructed to serve as a useful tool for exploring the population dynamics and behavioral effects of interacting species.

Our introduction of emotion to a predator-prey model has shown multiple biases in the population dynamics: increased food consumption; reduced predation; and increased population sizes. These biases are not guaranteed by the introduction of emotion as the dynamics of the system links

system attributes, i.e. too many rabbits can lead to overcrowding, decreased growth, and increased starvation. Nevertheless, the introduction of emotion to a species generally increases its population size, and it is to the species advantage to use emotions if it does not know if the competing species is also going to use emotion. Further work is required to understand the cooperative mode in the context of conspecific emotional communication; however, in our results we see favor towards neither species using emotion as the cooperative mode. This suggests that the acquisition of emotion may be an evolutionary result of competitive species interactions.

Communicated emotion can play a role in collective behavior for both predators and prey. Interestingly, the communication paradigm best for one species may be the worst for the other species. Rabbit emotions are more strongly affected by older information, whereas fox emotion is not. The fox population benefits when both populations use emotion if stimergetic communication is used with a low geometric decay. Thus, the predator is able to act collectively with a trail of information when it remains for a longer period of time. It will be interesting in the future to develop a system in which each species can evolve its communication strategies to see how the population dynamics and collective behavior are both affected by this additional dynamic.

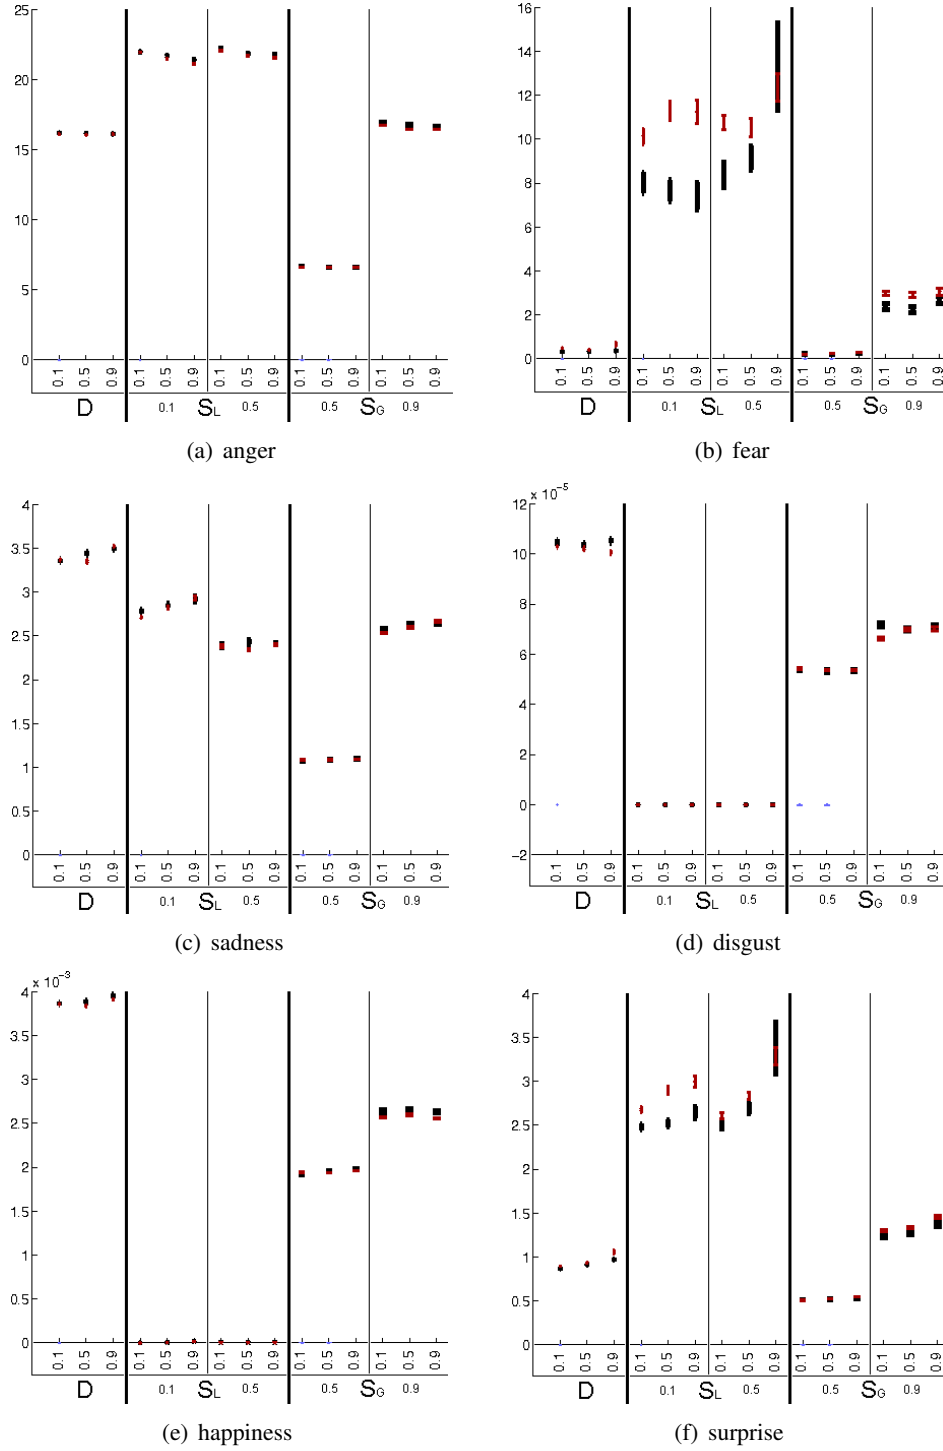


Figure 5.12. Average individual fox emotion for each parameter combination. The x-axis shows the communication coefficients (0.1,0.5,0.9) and decay values for each type of communication (N: no communication; D: direct; S: stigmergic (L=linear, G=geometric)). Black (darkest) represents when both species use emotion, Red (second darkest) represents when only foxes use emotion, and Blue (lightest) represents when only rabbits use emotion.

CHAPTER 6

EMOTIONS FOR AGENT COORDINATION

6.1 Introduction

Real-Time Artificial Intelligence has been investigated for over a decade [100]. A system is considered to be a Real-time AI system if it is able to make decisions within a guaranteed response time and thus meet domain deadlines. These systems face many challenges, including working with partial information, choosing the most crucial action if there are multiple scenarios to react to, and working continuously for an extended period of time without failure. These systems are usually created as expert systems, as they are used for a specific domain. However, they should be able to handle a wide variety of scenarios that may occur, not just specific test scenarios. Results must also be returned in a timely manner [100].

Real-Time Strategy (RTS) is an offshoot of general purpose real-time AI. RTS refers specifically to systems where the primary purpose is to create strategy, usually in a competitive atmosphere. For instance, military training on how to engage the enemy done via simulation is a RTS system. Only training with a computer strategy aspect is considered RTS, since it is not a RTS system if only the human controls strategy. Currently the military uses simulations heavily for training, and therefore it is crucial for them that these systems advance [68].

Although they may at first seem unrelated, emotions can play a large part in strategy especially when time is limited. Emotions are believed to improve our response time, increase our memory capacity, and provide quick communication [124]. We are able to notice things that we fear quicker than things we enjoy or are indifferent about, showing fear to be crucial to our response time. Remembering an emotion may enable a memory to be more useful for us later, as we can react to the emotion of the experience without needing to remember all of its details. Emotions help us convey our experience to another person; for instance, they will realize danger quicker from noticing our fear than by hearing our explanation. Thus, we propose to include emotion for collaboration between agents within a real-time strategy game. We provide each agent with simple emotions for

use in decision making and the ability to communicate them with their neighbors, similar to the emotions described in the Predator Prey system in Chapter 5.

Our system utilizes a current RTS gaming engine called Globulation that includes computer players that they call “AIs.” These are not to be confused with any specific algorithms in the field of AI, or with the field itself; however, we will retain this terminology to be consistent with the system. These AIs determine where its agents move, what they do, and when to create more of them; the same actions controlled by a human player. We provide computational emotions for these agents, and determine how those emotions affect the game play. We anticipate that emotions will enhance their ability to react to their environment and influence other agents, thus increasing the performance of the AI. One of our main contributions is the creation of an Emotion Sharing Map (ESM) that enables units to communicate their emotions in a way similar to the stigmergic communication discussed in Chapter 1. This Emotion Sharing Map saves the emotion of units and diffuses it for a period of time, enabling other units to feel the emotion of their peers. If the emotions are designed to be reactive to the environment, this map would enable agents to lay a trail for moving to or from specific types of areas without the need for either detailed or direct communication. This indirect communication between a single player’s agents allows for emergent behavior, where agents are able to work together just from following simple rules defining how their movement is modified by emotion.

In this chapter we will discuss related work in real-time strategy games, describe the system and the Emotion Sharing Map, show results, and then conclude.

6.2 Related Work

Real-time systems provide many new and difficult challenges for computer science. For example, a model of ship damage control has been created that relies on real time decision making [20]. This model determines the best course of action given the state of the ship and its many control systems. Tested in a simulation environment against actual Navy captains, the model vastly outperformed the humans. This example shows that Real-Time AI can even be valuable in situations where humans are already available to perform the task [20].

Also, many popular video games such as Starcraft and Warcraft incorporate Real-Time Strategy (RTS) if at least one team is computer controlled. These games all simulate war among multiple

players in which all but at least one player may be computer controlled. Although advances may be made in the AI of these systems, they do not seem to influence the military training development. However, many groups are working to combine the two groups so that meaningful work can be done to advance both fields at once [68, 23]. Ideally, the creation of war-related video games will be able to influence the military training simulations in years to come [22].

Real-time strategy games can involve many different fundamental AI issues. For instance, game AI is closely related to adversarial real-time planning, decision making under uncertainty, opponent modeling, spatial and temporal reasoning, resource management, collaboration, and path finding [23]. One system that is working to improve gaming in all of these aspects is ORTS [21]. This system is an open source game that is utilized in a competition each summer to encourage AI experts to test their skills and create software with a usable combination of solutions. Although we will use a similar system called “Globulation,” our enhancements could also be applied to ORTS.

Another way to create an RTS game is by controlling characters in games such as Quake. Laird et al. creates bots that can strategize through first person shooter games to beat human players [81]. They create their bots using real-time AI algorithms, giving them the ability to anticipate another player’s action, make smart decisions on where to go, and make smart decisions on what actions to take. This type of strategy is different from the type of strategy we will investigate, as it is only a single entity moving in a world against other similar entities [81].

Although there are currently no RTS systems that incorporate emotions that we are aware of, other software systems do exist with them. For instance, the digital life simulation game, the Sims, includes emotions. These emotions control the behavior of in-game agents; an example being that an unhappy agent is less likely to obey the commands of the controlling player. Many other examples of emotions being used in computer systems relate to the fields of human-computer interactions (HCI). A great amount of work has been done on improving a computer’s ability to detect a user’s emotions, and then using that information to change its interaction with the user. Much of this work is in the affective computing field [116, 1], and tends to relate to voice and facial recognition. A RTS system used for training can benefit from this work, but it is beyond our current scope.

6.3 The System

We modify the open source RTS platform Globulation (<http://www.globulation2.org>) by adding two emotions to the agents, modifying agent behavior based on their emotions, and sharing emotions among agents in a way similar to the communication paradigms already discussed in both the cancer chapter and the predator prey chapter through an Emotion Sharing Map. First we discuss the Globulation system, then the agents/units within Globulation, and finally the emotions and how they are utilized by agents.

6.3.1 Globulation

Globulation is a multi-player strategy war game where players compete for resources and territory, and the characters can be completely controlled by an AI. A player loses if all of their agents are destroyed.

Globulation has multiple AIs that can be chosen to act as an independent player in the game. The AI will control the actions of its assigned player so that no human intervention is needed. An AI defines a specific strategy, and will not only make overall player choices but will also give each agent its own set of decision processes. There are many different AIs available for Globulation, each with a different focus, level of detail, and success rate. The AI we will test against is named “Nicowar” and had the highest success rate of the AIs in our initial tests.

We define emotions as part of each agent’s individual decisions such that they can be ported to any of the AIs that already exist for the game. Thus, the decision processes for agents will be a combination of a previously created AI and our emotion and ESM combination. When defining our computational emotions we examined the deficiencies of the Nicowar AI. Although it is the most human-competitive AI in the game, its flaws include bottlenecks with path finding when dealing with a large numbers of agents, avoiding enemy agents (defensive agents), and finding enemy agents (offensive agents). We will seek to address all of these flaws with the emotions and ESM.

6.3.1.1 Agent Types

Each player in the game has their own agents that can be controlled by the overarching AI. These agents include warriors, workers, and explorers. Each agent has a numerical amount of health (HP) that can decrease if it is injured or increase if it is healed, with zero HP representing death. All

agents are capable of movement in 2-dimensional space within the map boundaries. They will make decisions on what actions to perform based on what they encounter as they move through the map. Our emotions will affect each agent's own decisions.

Each agent type has its own purpose in the game. The workers exist to gather resources needed for the player to build buildings, create more agents, and feed the current agents. The workers must coordinate so that they do not all approach the same resources at the same time. The warriors defend the player's buildings, and attack the opponent's buildings and agents. They must coordinate for both of these actions. The explorers will wander the map to determine the locations of enemies and resources, as these are not initially known. Thus, all agents are interdependent on the other agents for survival and a chance at winning the game. Emotions are given to workers and warriors, but not to explorers since they work independently and separately.

6.3.2 Emotions

6.3.2.1 Types of Emotions Modeled

We chose to model two different negative emotions which will be the same in each unit type, although each unit type will be affected differently by their emotions. The first emotion that we model is Fear, which is designed to keep agents alive and help coordinate warrior defense and offense. Fear is increased when a unit is attacked by an enemy unit, a unit is very damaged and close to death, or the player is running low on resources. The second emotions is Frustration, which is designed to combat path finding problems. Frustration is increased when a unit is unable to perform the task allotted to it or the unit has been on the same task for a significant amount of time.




Figure 6.1. The plane representing the range of a unit's emotions and 4 possible emotional stages: the origin is no fear or frustration, representing contentment; point 1 represents a unit with little frustration but high fear; point 2 is a unit with low fear and medium frustration; and point 3 represents a unit with high frustration and high fear.




Figure 6.2. A series of images demonstrating a player's Emotion Sharing Map changing over time. Each of the 6 images represents the entire game environment. Images are taken every 4,000 time steps. Frustration is shown in yellow(middle shade of gray), Fear is shown in red (darker shade), the overlap of the two emotions is green (lightest shade), and the lack of shared emotion is black. Images are organized chronologically from left-to-right and top-to-bottom.

Technically the lack of these two emotions also constitutes an emotion: contentment. For instance if there is little or no fear the unit feels content as the world seems safe. Also, if the unit has little or no frustration then it is content because everything is working well. Although units do not make decisions based on the combination of their 2 negative emotions, their emotional state at any time can be represented by a point on a plane with fear as the y-axis and frustration as the x-axis. A lack of emotion corresponds to contentment, as seen at the origin in Figure 6.1. However, without the the Emotion Sharing Map explained below, emotions would be entirely internal and not shared.

6.3.2.2 Emotion Sharing Map

For emotions to be most effective there must be a mechanism for agents to infer each other's emotions, as was seen in the predator/prey emotion chapter. For humans, emotions are exceptionally useful as a way to communicate. An agent's emotion is therefore influenced by the emotions of other agents under the same player via an Emotion Sharing Map. Agent emotions cannot be interpreted or felt by an opponent's units.

Each agent maintains a set of individual emotions that are modified based on the agent's actions and its reactions to its experiences. Each emotion is a continuum that we will assume is in some positive real valued range, with no loss of generality. The baseline for each agent is to have a value of zero for each emotion. Over time, any emotion increase will subsequently decrease until it reaches this baseline or a new experience replenishes it.

We define an Emotion Sharing Map (ESM) such that at each time step, an agent's internal emotions will be saved to the map. An example Emotion Sharing Map changing over time can be seen in Fig. 6.2. Each emotions is stored separately on the ESM. The ESM affects an agent's emotions and is updated by every agent's emotions at every time step. This frequency is to ensure that an agent has all information that may be vital to its decision making. The agent's emotion is added to the emotion on that square, and is immediately diffused out to the adjacent sets of squares within a specified grid square distance. A diffusion radius of 2 is shown in Fig. 6.3(a), assuming an agent is reacting to an experience with value 10. Given a *max_radius* defining the furthest distance an emotion is diffused and a *value* of the current emotion, the emotion value that will be saved on the map at a location that is *dist* away from the original point of the event is seen in Equation 6.1.

$$Map(dist) = \begin{cases} value & \text{if } dist = 0 \\ \frac{value - \frac{value}{max_radius}}{dist} & \text{otherwise} \end{cases} \quad (6.1)$$

Both the agent internal emotion values and the emotion values on the map decrease linearly over time. At each time step, the current emotions will decrease as shown in Fig. 6.3(b), and then any new emotions will be added. This type of emotional communication is a combination of the wide diffusion seen in the cancer chapters, and the continuous stigmergic style of communication seen in the predator-prey chapter. Emotions are diffused in all directions as the PLEASE DIE and I'M DYING signals were diffused in the cancer and HADES work, but are not reliant on specific events as were those signals. Instead they are saved at every time step as was seen in the predator-prey scenario.

Each agent can access a gradient of the map, and is affected by this gradient for each decision. The agent's own emotions are affected by the map such that a small percentage of each of its emotions is derived from the emotions on the map from the end of the previous time step, as defined in the next section. The map therefore allows agents to communicate indirectly, since the emotion values held on the map are due to another agent's recent experiences. Therefore, if an agent recently encountered a problem in a particular location, all close by agents will be aware due to the Emotion Sharing Map. Also, any other agents that come to the area within a short time span will be aware as well. The Emotion Sharing Map is therefore providing a mechanism for collaboration.

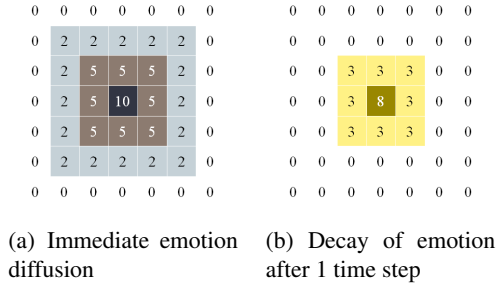


Figure 6.3. Approximate diffusion concept. The map in 6.3(a) depicts the values in the squares under and surrounding an agent that just experienced an event that resulted in a total emotion value of 10. If emotions decay linearly by 2, the map in 6.3(b) depicts the values in those same squares after a single time step before agents send their emotions to the map again.

6.3.2.3 Agents using Emotions

Each emotion affects units in ways related to five of Ekman’s seven characteristics of emotion: Quick onset, automatic appraisal, commonalities in antecedent events, brief duration, and unbidden occurrence [44]. Emotions occur based on events in an agent’s neighborhood immediately when that event occurs. The agent does not have time to decide that its surroundings are a problem, but instead there is a quick onset due to automatic appraisal of the situation. For all agents of a particular type the same antecedent event types will cause the same amount of the same emotion. Emotions are brief unless the same event continues to occur, in which case the emotion will continue to build at a slow rate. Emotions are not consciously caused, as only outside events or the sharing of emotions from another unit can cause them. The two characteristics that we do not relate to do not apply to our situation (presence in other primates, distinctive physiology) [44]. Actions that are taken due to an emotion are however decided upon only once the emotion reaches a specified threshold. Once that threshold is reached then the unit acts according to both its current situation and the fact that the particular emotion is strong.

An emotion’s effect on an agent is homogeneous throughout that agent type, although it differs between agent types. The effect of emotion is based on the idea of approach vs. withdraw [37]. In this theory, an emotion will elicit one of two responses: approaching toward the stimuli, or withdrawing from it. Emotion values are only incorporated into deciding an agent’s actions when the emotion value reaches a specified threshold. Each agent has two sets of controls: the built-in decision controls, and the emotion-based decision controls. Both sets of controls are potentially

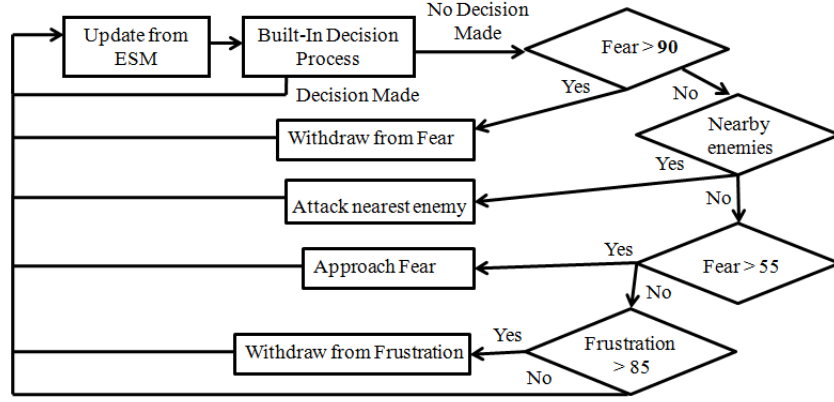


Figure 6.4. The decision tree for a warrior at each time step.

used at each time step, as can be seen in the warrior decision tree (Fig. 6.4). All updates from the ESM occur at the beginning of each time tick, and all emotions exist on a scale of 0 to 100.

Fear is affected by two factors: medical condition and surrounding enemy agents. If a worker's Fear is higher than 75 it will move in a direction toward less Fear until its Fear falls below that threshold, in an effort to save its own life. A warrior, however, will advance toward the source of Fear if their own fear is greater than 55. This reaction will cause a warrior to move toward nearby enemy agents and attack. However, if a warrior's Fear level rises higher than 90 it will retreat, improving on its ability to survive. The value for Fear (Υ) of an agent at time t in location λ if it is surrounded by ϕ enemies is shown in Eq. 6.2 where $Map(\Upsilon, \lambda)$ refers to the value of Fear on the ESM in location λ and ω is 1 if the agent is damaged and 0 otherwise.

$$\Upsilon(t) = 0.8 \cdot \Upsilon(t-1) + 10\phi + \omega + 0.1 \cdot Map(\Upsilon, \lambda) \quad (6.2)$$

The agent reactions are similar for Frustration. If a worker has a Frustration level over 85 it will look elsewhere for work, which will usually involve looking for resources to gather. If the worker is already in a location with resources but still has high Frustration, it is likely due to a large number of workers gathered who are causing a bottleneck for retrieving resources. If a warrior has Frustration it will explore to look for enemies or will wander around acting as a lookout, as Frustration is likely a result of no danger in its current location. Frustration directly combats the AI's problem of

failed path finding. Thus Frustration can create a more efficient resource gathering mechanism for workers, and a higher likelihood of encountering enemies for warriors.

Frustration is increased in a particular agent by one tenth of the amount of time spent continuously performing the same task. This increase of Frustration allows agents stuck in a location to free themselves by moving away from the Frustration gradient. An agent's value for Frustration (Ω) at time t can thus be similarly set as seen in Equation 6.3 if χ is a binary number that is 1 if `actionTickTimer` represents the time the agent has been doing the same action, ($\text{actionTickTimer} > 50$), and ($\text{actionTickTimer} \% 10 = 0$).

$$\Omega(t) = 0.8 \cdot \Omega(t - 1) + \chi + 0.1 \cdot \text{Map}(\Omega, \lambda) \quad (6.3)$$

Although we discuss emotions as causing workers to avoid death and find resources, and warriors to find enemies and avoid death, this behavior is not hard-coded into the system. Instead, this behavior emerges from the agents following the simple rules based on how to react to emotions and share emotions through the Emotion Sharing Map.

6.4 Experimental Design

Simulations were run with version 0.9.1 of Globulation 2 on the map Muka, which is a one player versus one player map. Each player has all necessary resources contained within a region that is connected to the opponent via two land bridges (at the top and bottom). The map wraps from right to left, creating a land bridge from the left side to the right side of each player's region. Both players also have an additional smaller peninsula containing resources. The map is essentially symmetric, to make each player's starting situation close to identical.

We test the Nicowar AI against itself both with and without the ESM to determine if the AI was improved by using emotion. For simplicity of explanation, we will call Nicowar using emotions "NicowarESM", and Nicowar without emotions "Nicowar." NicowarESM will always be player 1 and Nicowar will always be player 2. We test two AIs against each other instead of against humans to reduce the error that may occur from a human changing their strategy over time.

We test variations on the constant diffusion radius for both Fear and Frustration (Table 6.1) to determine how far the messages must be communicated to be effective in positively influencing

Fear	1	2	2	2	3
Frustration	2	1	2	3	2

Table 6.1. Diffusion radii tested for fear and frustration.

agent behavior. Each set of parameters is tested eight times. We compare player 1 to player 2 in each scenario to ensure that our results are not biased due to starting location on the map, as the Nicowar vs Nicowar results are also calculated this way.

We use two statistics for analyzing success: percent of games won, and average health (HP) for each player’s agents and buildings. A high HP per agent ratio can signify that either the player has a high number of agents in various stages of health, or that all agents have high HP. High HP can signify more powerful warriors as well. Since all of these scenarios can represent a successful game, they also imply good performance.

Average health for each player is determined by examining the hit points per agent and per building for each player over the course of the game. Since each of these games is two AIs playing against each other, we can take the difference of their HP ratios at each time step and then average them. This average represents how much better the HP/agent ratio for player 1 is over player 2 for the duration of the game. A game is won when all of the other player’s buildings and agents are destroyed.

6.5 Results

Results are presented with varying diffusion radiuses of the form “diffusion radius of Fear, diffusion radius of Frustration,” for instance ESM(1,2) for diffusion radius of 1 for Fear and 2 for Frustration. The radius values can be seen in Table 6.1.

As can be seen in Fig. 6.5, the percent of wins generally increases when the ESM is utilized to facilitate collaboration and communication. For NicowarESM(1,2) and NicowarESM(2,2) the win percentage is more than double the win percentage of the baseline. The baseline has a win percentage less than 50% as the map appears slightly biased against Player 1. However, since we have used Player 1 as the player with the ESM, switching players would not decrease our results.




Figure 6.5. Percent of the eight games won by NicowarESM (Player 1) against Nicowar (Player 2) with varying diffusion radii (as labeled). The dashed line shows the baseline, i.e. the percent of wins by Nicowar when playing itself without emotion.

A carefully chosen diffusion radius (the distance away the emotion is shared) can greatly improve the number of wins. From the wins alone, it appears that a diffusion radius of 1 for Fear, and a diffusion radius of 2 for Frustration is ideal. Since battles may include a large number of agents, a lower diffusion radius for Fear is probably desired as a way to keep the map from becoming too overrun with emotions. Since Frustration is generally only caused by agents from the same player, a higher diffusion radius allows trapped agents to keep others from further causing problems by blocking them, thus freeing themselves sooner.

All combinations tested except diffusion radius for Fear of 3 with a diffusion radius for Frustration of 2 improve on the baseline. Once the diffusion radius increases too high the ESM most likely becomes harder to use for navigation due to a high number of emotions mingling such that it is difficult to determine which direction is the correct way to turn.

The difference of average HP for agents and buildings between players at least doubles over the Nicowar baseline for all NicowarESM results (Fig. 6.6(a)). Although NicowarESM(3,2) does not increase in percentage of wins, it does increase significantly in the average HP of agents and buildings. This may be caused by NicowarESM keeping its agents to safe and healthy that it in turn prevents its agents from adequately attacking the enemy, causing them to eventually be destroyed. NicowarESM(2,2) has a similar situation in that although it is the second best by number of wins it has the highest overall HP of any scenarios tested.




Figure 6.6. Difference of average hp per agents and buildings from Nicowar (player 2) versus NicowarESM (player 1) with varying diffusion radii (as labeled). (a) When examining all experiments, we see that the most significant increase in HP is with diffusion radius of (2,2). (b) If we only examine cases where player 1 won the game, not only do most average HPs increase, but also the diffusion radius of (3,2) increases higher above Nicowar’s HP.

Either when these AIs win they have a much higher HP at the end, or they maintain a higher HP throughout most games until they lose. In Figure 6.6(b) we examine the average HP only in the cases where player 1 won the game. NicowarESM(3,2) has the highest average HP in this case, so much of its variation in Figure 6.6(a) is probably due to high average HP when it wins, but not particularly high HP during games it loses. NicowarESM(2,2)’s average HP does not increase as much when only considering games it wins, so it is likely that this set of diffusion radii gives an advantage in HP both when it wins and loses.

If we consider all cases instead of only when the AI wins the game, then from the average health per agent and building statistics the best results are a diffusion radius for Fear and Frustration of 2, although a Fear radius of 3 with a Frustration radius of 2 is a close second. Taking percent of wins into account as well shows an overall winner when Fear is 1 and Frustration is 2, since the number of wins is increased more and the agents still maintain higher health and total numbers.

We thus show evidence toward the ESM enhancing the coordination of agents in Globulation. Different diffusion radii can be used to either significantly increase the number or wins, or to significantly increase the health of the player’s agents and buildings. The use of computational emotion increases the survivability of the player, at least when that player is the only one using emotions.

This result mimics what was seen in the Predator Prey modeling, where a similar emotion communication increased the survivability of both predator and prey when they were the only species using emotion. Since these real-time strategy games mimic population dynamics in many ways,

this result should be expected. However, it is interesting to note that just a sharing of emotion with neighbors on a short time scale can cause a global effect of increased self-organization and coordination in a group of agents that is generally not collaborating in any real sense.

6.6 Conclusions

In this chapter we have developed computational emotions for real-time strategic games, and tested these emotions in a game called Globulation. Globulation has autonomous computer players they call “AI”s, which control a player’s agents. We provide two emotions that influence an agent’s movement in the game: Fear and Frustration. We provide an Emotion Sharing Map that agents use to communicate simple information about their recent experiences. This information remains on the ESM until it decays to zero, influencing the internal emotions of nearby agents. Through communication on the Emotion Sharing Map, agents are able to collaborate such that Frustration improves path finding issues, and Fear helps agents survive longer and helps warriors find battles.

We show improvement to Globulation’s top AI “Nicowar” by adding emotions and the ESM. Our results show that the ESM is an effective way to allow agents to communicate their emotions with neighbors to aid in collaboration. It does not require direct communication but is more reminiscent of swarm communication. Emotions are diffused immediately within a specific radius, and at each time step an agent’s internal emotions will be affected by the ESM emotions at their current location. This use of emotions is very similar to what was discussed in the previous predator prey chapter, however emotions are diffused further and faster in Globulation. This modification on the predator prey emotions increases the speed and accuracy of information, and since this is an entirely computational system there are no odd effects as may occur if using a similar communication scheme in a biological model.

Nicowar AI wins twice as many games when it uses emotions, and retains more agents of better health overall. In all but one set of diffusion radii tested Nicowar wins more games when it uses emotion, and with all sets of diffusion radii its average HP increases when it uses emotions. There are a variety of diffusion radii that may be chosen for this system, all of which will improve performance by some degree. This emotion technique is thus potentially applicable to other real-time strategy games, or real-time collaboration systems. Both emotions were designed specifically to counteract problems in the Nicowar AI, which may imply that emotions designed for a specific

system will improve that system's functionality. It may be possible to create computational emotions that will work in general for all strategies, but further investigation is necessary.

CHAPTER 7

CONCLUSIONS AND FUTURE DIRECTIONS

7.1 Conclusions

In this dissertation I presented two approaches to uniting biological complex systems with computer science: through modeling, and through inspiration. I showed new approaches to modeling cancer and predator prey systems, as well as using computational models of emotion to improve agent communication and predator-prey decision making. Additionally, I examined how the communication model from cancer can be re-applied to multi-agent fault tolerance. Overall, this dissertation supports six claims:

- Claim 1** *Varying spatial regulations within a cellular system leads to significantly different number of predicted cells in a model of cancer cell growth.*
- Claim 2** *Intercellular messaging among cells based on neighbor death and spatial impingement can be used to encourage death of surrounding cells such that primarily cancer cells are killed and healthy tissue cells survive.*
- Claim 3** *The intercellular messaging investigated for cancer removal can improve multi-agent system fault tolerance by allowing agents to use only local information and collaboration to remove faulty agents.*
- Claim 4** *Computational emotions provide a framework for modeling predator-prey dynamics that will provide different modeling behavior than a traditional cellular automata model.*
- Claim 5** *Computational emotions can be used to improve the performance of a computer player in a real-time strategy game.*
- Claim 6** *Stigmergic communication can be utilized to improve collaboration in both modeling interdisciplinary problems and when designing computational systems.*

In all of these situations there is some form of emergent behavior, as simple rules cause complex behavior. Often this is accomplished by communication of information about the world to neighbors. This form of communication is a modification of Stigmergy. In both the cancer agent-based model and the multi-agent fault tolerance system HADES, death request signals are sent in response to specific local events related to pushing either a cell or agent out of its current location. These signals diffuse in all directions for a specific radius, and are felt strongest by the closest neighbors of the sender. Eventually, these signals are able to induce death/removal of the recipients and can rid the system of cancer cells or irregular agents, respectively. We show that for cancer this type of communication mimics existing mechanisms, and may explain the body's ability to fight cancer naturally. For multi-agent systems we show that this type of communication may be used to provide fault tolerance, by combining two fault tolerance techniques: Survivalist and Citizen. Anonymous agent communication enables removal of malfunctioning agents, requiring only knowledge of local phenomena.

Two variations on this type of communication are used to develop a new approach to modeling predator-prey dynamics in cellular automata. The internal emotional state of each entity within the system is shared with neighbors. That shared information is then incorporated into the neighbor's own internal emotional state to influence its future decisions. We test two approaches to communication between conspecifics: stigmergic communication in which a decaying trail is left as the entity moves, and direct communication in which only current neighbors are told about an entity's emotion and thus there is no diffusion. We find that both predator and prey benefit from the use of emotions in their decision making, and that the form of communication significantly affects both populations.

A similar form of computational emotions is then applied to agents in a real-time strategic game. These emotions influence an agent's behavior to combat issues in the controller related to pathfinding and fighting enemies. In this case we use communication that is a combination of the approaches in the other systems. Messages are diffused in all directions as in the cancer approach, but linger and decay over time as in the predator-prey approach. We find that the use of communicated emotions increases collaboration among agents, and enables the computer controlled player to win the game almost twice as frequently as it did without the use of emotions.

7.2 Future Directions

This dissertation provides many future directions for each of the systems tested, as well as for stigmergic communication.

7.2.1 Cancer

Cancer is a very complicated disease that we do not yet completely understand. There are many modeling opportunities, both building upon the models presented here as well as developing new models with similar approaches. One possibility is to examine how different therapies work against cancer. There are already models specifically for therapies such as radiation therapy and chemotherapy. It would be useful to compare how these models work with and without the communication protocols presented in this dissertation. If these communication protocols are indeed describing mechanisms that occur in nature, they may also influence how well these therapies work. Thus, a future possibility is to modify the model to incorporate these therapies and determine what version of the model gives the most accurate results. This would also provide another way to analyze the realism of the communication protocols.

Another question in cancer that could be examined using a similar type of agent-based model is the role of cancer stem cells. It is generally agreed upon that at least some cancer cells act like stem cells in their ability to proliferate frequently. However, there is a debate on whether stem cells become cancer cells or if cancer cells de-differentiate into stem cells. Using a similar agent-based model approach, we can examine which scenario leads to the correct rate of cancer formation, time to cancer, and cancer growth shape.

These are only two potential directions for working on cancer, but there are many other questions, and doubtless new questions will be generated from biological studies that could also benefit from a computational approach.

7.2.2 Multi-agent Fault Tolerance

We present a set of communication protocols to remove malfunctioning agents in a generic multi-agent system. We treat this system as an abstraction to other systems, in which our defined properties could be slightly re-interpreted to apply the communication protocols to a specific system. The primary next step is to take standard systems such as sensor networks, distributed software

systems, or others, and more fully develop those analogies to test the communication protocols on specific failures in these systems. This work will involve testing variations on the communication protocol, as well as testing it on variations of the network and system structure. The system presented in this dissertation shows promise, and thus is likely to work in these other systems as well. This future work may also result in additional versions of the protocol that can then be applied to specific classes of problems as well.

7.2.3 Computational Emotions for Predator Prey Modeling

The main direction for this work is examining it for modeling of specific biological predator-prey systems. It is not expected that an emotional cellular automata approach would encompass the dynamics of all predator-prey systems, and thus the first problem is to determine exactly which systems could benefit from this modeling technique. The next problem would be determining if this technique can better predict changes in that biological system by testing it on population variance data.

There are also many modeling approaches that could be further tested in this type of system as well. The way that rabbits and foxes made decisions using their emotions was fixed throughout experiments, but may show interesting changes if modified. There was also no variation on how emotions were affected by the environment, which may change rabbit and fox behavior significantly.

7.2.4 Computational Emotions for Real-time Games

The emotions worked well in solving the problems in the computer controlled player. These emotions could be developed into an overall architecture that could be applied to many different players (using different strategies) in both this real-time strategy game as well as other similar games. This would involve trying different versions of computational emotions and determining when they do and do not work well. Other types of communication could also be tested to determine if this approach is the only approach that works, or if there are other possibilities. We could also enable learning in the agents, where a constantly high emotion in one region is remembered and over time they learn to avoid (or approach) that region.

BIBLIOGRAPHY

- [1] *Affective Computing and Intelligent Interaction* (2007), vol. 4738/2007, Springer.
- [2] Abbott, Robert. Cancersim: A computer-based simulation of hanahan and weinberg's hallmarks of cancer. Master's thesis, University of New Mexico, 2002.
- [3] Adamatzky, A. Affectons: automata models of emotional interactions. *Applied Mathematics and Computation* 146 (2003), 579–594.
- [4] Adamatzky, A., and Grube, M. On localizations in minimal cellular automata model of two-species mutualism. *Int. J. Bifurcation and Chaos* (2009).
- [5] Allman, Elizabeth S., and Rhodes, John A. *Mathematical Models in Biology*. Cambridge University Press, 2004.
- [6] Allsopp, R C, Vaziri, H, Patterson, C, Goldstein, S, Younglai, E V, Futcher, A B, Greider, C W, and Harley, C B. Telomere length predicts replicative capacity of human fibroblasts. *Proceedings of the National Academy of Sciences of the United States of America* 89, 21 (1992), 10114–10118.
- [7] Anderson, Alexander R.A., Weaver, Alissa M., Cummings, Peter T., and Quaranta, Vito. Tumor morphology and phenotypic evolution driven by selective pressure from the microenvironment. *Cell* 127, 5 (2006), 905 – 915.
- [8] Anderson, P.W. More is different. *Science* 177, 4047 (August 1972).
- [9] Apetoh L, Tesinire A, Ghiringhelli F Kroemer G Zitvogel L. Molecular interactions between dying tumor cells and the innate immune system determine the efficacy of conventional anti-cancer therapies. *Cancer Res* 68 (2008), 4026–4030.
- [10] Ashkenazi, Avi, Holland, Pamela, and Eckhardt, S. Gail. Ligand-Based Targeting of Apoptosis in Cancer: The Potential of Recombinant Human Apoptosis Ligand 2/Tumor Necrosis Factor-Related Apoptosis-Inducing Ligand (rhApo2L/TRAIL). *J Clin Oncol* 26, 21 (2008), 3621–3630.
- [11] Barkai, Naama, and Shilo, Ben-Zion. Variability and robustness in biomolecular systems. *Molecular Cell* 28, 5 (2007), 755 – 760.
- [12] Bauer, Amy L., Jackson, Trachette L., and Jiang, Yi. A cell-based model exhibiting branching and anastomosis during tumor-induced angiogenesis. *Biophysical Journal* 92, 9 (2007), 3105–3121.
- [13] Bechara, A., Damasio, H., and Damasio, A. R. Emotion, decision making and the orbitofrontal cortex. *Cerebral Cortex* 10 (2000), 295–307.

- [14] Bell, A.V., Rader, R.B., Peck, S.L., and Sih, A. The positive effects of negative interactions: Can avoidance of competitors or predators increase resource sampling by prey? *Theoretical Population Biology* 76 (2009).
- [15] Blanchard, D.C., Griebel, G., and Blanchard, R.J. Conditioning and residual emotional-ity effects of predator stimuli: some reflections on stress and emotion. *Progress in Neuro-Psychopharmacology and Biological Psychiatry* 27 (2003), 1177–1185.
- [16] Bokareva, Tatiana, Bulusu, Nirupama, and Jha, Sanjay. Sasha: Toward a self-healing hybrid sensor network architecture. *Embedded Networked Sensors* (2005), 30–31.
- [17] Bonabeau, Eric, Dorigo, Marco, and Theraulaz, Guy. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, 1999.
- [18] Bonabeau, Eric, Dorigo, Marco, and Theraulaz, Guy. Inspiration for optimization from social insect behaviour. *Nature* 406, 6791 (2000), 39–42.
- [19] Bru, Antonio, Albertos, Sonia, Subiza, Jose Luis, Garcia-Asenjo, Jose Lopez, and Bru, Isabel. The Universal Dynamics of Tumor Growth. *Biophys. J.* 85, 5 (2003), 2948–2961.
- [20] Bulitko, Vadim, and Wilkins, David. Real-time decision making for shipboard damage control. In *AAAI* (2001).
- [21] Buro, Michael. Orts: A hack-free rts game environment. In *Proceedings of the International Computers and Games Conference* (2002).
- [22] Buro, Michael, and Furtak, Timothy. Rts games as test-bed for real-time research. In *Workshop on Game AI, JCIS* (2003).
- [23] Buro, Michael, and Furtak, Timothy. Rts games and real-time ai research. In *Proceedings of the Behavior Representation in Modeling and Simulation Conference (BRIMS)* (2004).
- [24] C.Ferreira, S., Martins, M. L., and Vilela, M. J. Reaction-diffusion model for the growth of avascular tumor. *Phys. Rev. E* 65, 2 (Jan 2002), 021907.
- [25] Chaplain, Mark A. J. Mathematical modelling of angiogenesis. *Journal of Neuro-Oncology* 50 (2000), 37–51. 10.1023/A:1006446020377.
- [26] Chaplain, Mark A. J., McDougall, Steven R., and Anderson, Alexander. R. A. Mathematical modeling of tumor-induced angiogenesis. *Annual Review of Biomedical Engineering* 8, 1 (2006), 233–257.
- [27] Chen, Q., and Mynett, A.E. Effects of cell size and configuration in cellular automata based prey-predator modelling. *Simulation Modelling Practice and Theory* (2003).
- [28] Cheng, Jimming, Cheng, Winston, and Nagpal, Radhika. Robust and self-repairing formation control for swarms of mobile agents. In *Proceedings of the Twentieth National Conference on Artificial Intelligence* (Menlo Park, California, 2005), AAAI Press, pp. 59–64.
- [29] Coffey, Donald S. Self-organization, complexity and chaos: The new biology for medicine. *Nature Medicine* 4, 8 (August 1998), 882–885.
- [30] Couzin, Iain. Collective minds. *Nature* 445 (2007).

- [31] Couzin, Iain D. Collective cognition in animal groups. *Trends in Cognitive Sciences* 13 (2009), 36–43.
- [32] Couzin, Iain D., Krause, Jens, Franks, Nigel R., and Levin, Simon A. Effective leadership and decision-making in animal groups on the move. *Nature* 433 (2005), 513–516.
- [33] Coyle, E. A., Maguire, L. P., and McGinnity, T. M. Self-repair of embedded systems. *Engineering Applications of Artificial Intelligence* 17, 1 (2004), 1–9.
- [34] Curtis, V., Aunger, R., and Rabie, T. Evidence that disgust evolved to protect from risk of disease. In *Proceedings of the Royal Society of London, Series B: Biological Sciences* (2004), vol. 271, pp. 131–133.
- [35] Damasio, Antonio. *Looking for Spinoza: Joy, Sorrow, and the Feeling Brain*. Harcourt Books, 2003.
- [36] Damasio, A.R., Tranel, D., and Damasio, H. *Somatic markers and the guidance of behaviour: theory and preliminary testing*. New York: Oxford University Press, 1991, pp. 217–229.
- [37] Davidson, Richard J., Ekman, Paul, Saron, Clifford D., Senulis, Joseph A., and Friesen, Wallace V. Approach-withdrawal and cerebral asymmetry: Emotional expression and brain physiology i. *Journal of Personality and Social Psychology* 58, 2 (1990), 330–341.
- [38] Davis, D.N. Cellular automata, computational autonomy and emotion. In *International Conference on Computational Intelligence for Modelling, Control and Automata* (2001).
- [39] de Carvalho, K.C., and Tome, T. Self-organized patterns of coexistence out of a predator-prey cellular automaton. *International Journal of Modern Physics C* 17 (2006), 1647–1662.
- [40] Dewdney, A.K. Sharks and fish wage an ecological war on the toroidal planet wa-tor. *Scientific American* (1984).
- [41] DJ, McConkey. Therapy-induced apoptosis in primary tumors. *Adv Exp Med Biol* 608 (2007), 31–51.
- [42] Dong, C.Y., Long, J.T., Reiter, C.A., Staten, C., and Umbrasas, R. A cellular model for spatial population dynamics. *Computers & Graphics* 34 (2010), 176–181.
- [43] Dréau, Didier, Stanimirov, Dimitre, Carmichael, Ted, and Hadzikadic, Mirsad. An agent-based model of solid tumor progression. In *Proceedings of the 1st International Conference on Bioinformatics and Computational Biology* (Berlin, Heidelberg, 2009), BICoB '09, Springer-Verlag, pp. 187–198.
- [44] Ekman, Paul. *All Emotions are Basic*. Oxford University Press, 1994.
- [45] Ekman, Paul. *Basic Emotions*. Sussex, U.K.: John Wiley & Sons, Ltd., 1999.
- [46] Enderling, Heiko, Chaplain, Mark A.J., Anderson, Alexander R.A., and Vaidya, Jayant S. A mathematical model of breast cancer development, local treatment and recurrence. *Journal of Theoretical Biology* 246 (2007), 245–259.
- [47] Eng, C., Leone, G., Orloff, MS, and Ostrowski, M.

- [48] Farina, F., and Dennunzio, A. A predator-prey cellular automaton with parasitic interactions and environmental effects. *Fundam. Inf.* 83 (2008), 337–353.
- [49] Fedoruk, Alan, and Deters, Ralph. Improving fault-tolerance by replicating agents. In *AA-MAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems* (New York, NY, USA, 2002), ACM Press, pp. 737–744.
- [50] Francis, K., and Palsson, B. Effective intercellular communication distances are determined by the relative time constants for cyto/chemokine secretion and diffusion. *Proceedings of the National Academy of Sciences* 94 (November 1997), 12258–12262.
- [51] Frank, S. A., Iwasa, Y., and Nowak, M. A. Patterns of cell divisions and the risk of cancer. *Genetics* 163 (April 2003), 1527–1532.
- [52] Fulda, Simone. Inhibitor of apoptosis proteins as targets for anticancer therapy. *Expert Review of Anticancer Therapy* 7, 9 (2007), 1255–1264.
- [53] G, Abbott, S, Forent, and KJ, Pienta. Simulating the hallmarks of cancer. *Artificial Life* 12 (2006), 617–34.
- [54] Galam, Serge, and Radomski, Jan P. Cancerous tumor: The high frequency of a rare event. *Phys. Rev. E* 63, 5 (Apr 2001), 051907.
- [55] Gardner, Martin. The fantastic combinations of john conway’s new solitaire game ‘life’. *Scientific American* 223 (1970), 120–123.
- [56] George, Selvin, Evans, David, and Marchette, Steven. A biological programming model for self-healing. In *SSRS '03: Proceedings of the 2003 ACM workshop on Survivable and self-regenerative systems* (New York, NY, USA, 2003), ACM Press.
- [57] Gerisch, A., and Chaplain, M.A.J. Mathematical modelling of cancer cell invasion of tissue: Local and non-local models and the effect of adhesion. *Journal of Theoretical Biology* 250 (2008), 684–704.
- [58] Gerlee, P., and Anderson, A.R.A. An evolutionary hybrid cellular automaton model of solid tumour growth. *Journal of Theoretical Biology* 246 (2007), 583–603.
- [59] GJ, GJ Pettet, Please, CP, Tindall, MJ, and McElwain, DLS. The migration of cells in multi-cell tumor spheroids. *Bulletin of Math Bio* 63 (2001), 231–257.
- [60] Goldberg, David E. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Professional, Jan. 1989.
- [61] Griffith, S., Goldwater, D., and Jacobson, J. Self-replication from random parts. *Nature* 437 (2005), 636.
- [62] Gronewold, A., and Sonnenschein, M. Event-based modelling of ecological systems with asynchronous cellular automata. *Ecological Modelling* 108 (1998), 37–52.
- [63] Hackwood, Susan, and Beni, Gerardo. Self-organization of sensors for swarm intelligence. In *Proceedings of the IEEE International Conference on Robotics and Automation* (1992).
- [64] Hamscher, W., Console, L., and de Kleer, J. *Readings in Model-Based Diagnosis*. Morgan Kaufmann Publishers Inc., 1992.

- [65] Hanahan, D., and Weinberg, R. A. The hallmarks of cancer. *Cell* 100 (January 2000), 57–70.
- [66] Handl, Julia, and Meyer, Bernd. Ant-based and swarm-based clustering. *Swarm Intelligence* 1 (2007), 95–113.
- [67] Hawick, K.A., and Scogings, C.J. A minimal spatial cellular automata for hierarchical predator-prey simulation of food chains (tech. rep. cstn-040). Tech. rep., Computer Science, Massey University., 2009.
- [68] Herz, J.C., and Macedonia, M.R. Computer games and the military: Two views. Tech. rep., Center for Technology and National Security Policy, National Defense University, April 2002.
- [69] Hess, Corine J., Berkhof, Johannes, Denkers, Fedor, Ossenkoppele, Gert J., Schouten, Jan P., Oudejans, Joost J., Waisfisz, Quinten, and Schuurhuis, Gerrit J. Activated Intrinsic Apoptosis Pathway Is a Key Related Prognostic Parameter in Acute Myeloid Leukemia. *J Clin Oncol* 25, 10 (2007), 1209–1215.
- [70] Hinchey, M.G., Sterritt, R., and Rouff, C. Swarms and swarm intelligence. *Computer* 40, 4 (April 2007), 111–113.
- [71] Hogeweg, P. Cellular automata as a paradigm for ecological modeling. *Applied Mathematics and Computation* 27 (1988), 81–100.
- [72] Horling, B., Lesser, V., Vincent, R., Bazzan, A., and Xuan, P. Diagnosis as an integral part of multi-agent adaptability. *Proceedings of DARPA Information Survivability Conference and Exposition* (January 2000), 211 – 21.
- [73] Ingerson, T.E., and Buvel, R.L. Structure in asynchronous cellular automata. *Physica D: Nonlinear Phenomena* 10 (1984), 59–68.
- [74] JB, Swann, and MJ, Smyth. Immune surveillance of tumors. *Journal of Clinical Investigations* 117, 5 (2007), 1137–1146.
- [75] Johansson, A., and Sumpter, David J.T. From local interactions to population dynamics in site-based models of ecology. *Theoretical Population Biology* 64 (2003), 497–517.
- [76] Kamboj, Sachin. Analyzing the tradeoffs between breakup and cloning in the context of organizational self-design. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 2* (Richland, SC, 2009), AAMAS '09, International Foundation for Autonomous Agents and Multiagent Systems, pp. 829–836.
- [77] Kaplan, Daniel, and Glass, Leon. *Understanding Nonlinear Dynamics*. Springer-Verlag, 1995.
- [78] Khain, E., and Sander, L. M. Dynamics and pattern formation in invasive tumor growth. *Physical Review Letters* 96, 188103 (2006), 1–4.
- [79] Kitano, Hiroaki. Systems biology: A brief overview. *Science* 295, 5560 (2002), 1662–1664.
- [80] Klein, M., Rodriguez-Aguilar, J., and Dellarocas, C. Using domain-independent exception handling services to enable robust open multi-agent systems: The case of agent death. *Autonomous Agents and Multi-Agent Systems* 7 (2003), 179–189.

- [81] Laird, John. Using a computer game to develop advanced ai. *IEEE Computer* (2001).
- [82] Lansing, J. Stephen. Complex adaptive systems. *Annu. Rev. Anthropol.* 32 (2003), 183–204.
- [83] Lehman, Clarence L., and Tilman, David. *Competition in Spatial Habitats*. Princeton University Press, NJ, 1997, pp. 185–203.
- [84] LH, Abbott, and F, Michor. Mathematical models of targeted cancer therapy. *Br J Cancer* 95 (2006), 1136–1141.
- [85] Lindahl, Tomas, and Wood, Richard D. Quality control by dna repair. *Science* 286, 3 (December 1999).
- [86] Lotka, A.J. *Elements of Physical Biology*. Williams and Wilkins, 1925.
- [87] Low, A., Lang, P.J., Smith, J.C., and Bradley, M.M. Both predator and prey: Emotional arousal in threat and reward. *Psychological Science* 19 (2008), 865–873.
- [88] Mange, Daniel, Sipper, Moshe, Stauffer, Andre, and Tempesti, Gianluca. Toward self-repairing and self-replicating hardware: The embryonics approach. *EH 00* (2000), 205.
- [89] Mantzaris, NV, and Webb, S annd Othmer, HG. Mathematical modeling of tumor-induced angiogenesis. *Journal of Mathematical Biology*, 2 (2004), 111–87.
- [90] Marin, Olivier, Sens, Pierre, Briot, Jeanpierre, Guessoum, Zahia, and Cedex, Bp Le Havre. Towards adaptive fault tolerance for distributed multi-agent systems. In *Proceedings of European Research Seminar on Advances in Distributed Systems* (2001).
- [91] Markus, Mario, Böhm, Dominik, and Schmick, Malte. Simulation of vessel morphogenesis using cellular automata. *Mathematical Biosciences* 156, 1-2 (1999), 191–206.
- [92] Marusic, M., Bajzer, Z., Vuk-Pavlovic, S., and Freyer, J.P. Tumor growth in vivo and as multicellular spheroids compared by mathematical models. *Bulletin of Mathematical Biology* 56 (1994), 617–631.
- [93] Matsumoto, Makoto, and Nishimura, Takuji. Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Trans. Model. Comput. Simul.* 8, 1 (1998), 3–30.
- [94] McDougall, Steven R., Anderson, Alexander. R. A., and Chaplain, Mark A. J. Mathematical modelling of dynamic adaptive tumour-induced angiogenesis: Clinical implications and therapeutic targeting strategies. *Journal of Theoretical Biology* 241, 3 (2006), 564–589.
- [95] Meinhardt, H., and Gierer, A. Generation and regeneration of sequence of structures during morphogenesis. *Journal of Theoretical Biology* 85, 3 (August 1980), 429–50.
- [96] Michor, Franziska, Iwasa, Yoh, and Nowak, Martin A. Dynamics of cancer progression. *Nat Rev Cancer* 4 (2004), 197–206.
- [97] Miller, J. Evolving a self-repairing, self-regulating, french flag organism. In *Proceedings of GECCO* (2004).
- [98] Millonas, Mark M. Swarms, phase transitions, and collective intelligence. In *SFI Studies in the Sciences of Complexity* (1994), vol. 17, pp. 417–417.

- [99] Mitzenmacher, Michael. A brief history of generative models for power law and lognormal distributions. *Internet Mathematics* 1, 2, 226–251.
- [100] Musliner, David J., Hendler, James A., Agrawala, Ashok K., Durfee, Edmund H., Strosnider, Jay K., and Paul, C.j. The challenges of real-time ai. *Computer* 28, 1 (January 1995), 58–66.
- [101] N, Beerenwinkel, T, Antal, D, Dingli, A, Traulsen, and Kinzler KW, et al. Genetic progression and the waiting time to cancer. *PLoS Comput Biol* 3 (2007), 2239–2246.
- [102] Novark, G., Berger, E. D., and Zorn, B. G. Exterminator: Automatically correcting memory errors with high probability. In *Proceedings of the Conference on Programming Language Design and Implementation* (July 2007).
- [103] Nowak, Martin A. *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press, 2006.
- [104] Olsen, Megan, Harrington, Kyle, and Siegelmann, Hava. Emotions for strategic real-time systems. In *AAAI Emotion, Personality, and Social Behavior Technical Report (SS-08-04)* (March 2008), pp. 104–110.
- [105] Olsen, Megan, Harrington, Kyle, and Siegelmann, Hava. Conspecific emotional cooperation biases population dynamics: a cellular automata approach. *International Journal of Natural Computing Research* 1, 3 (2010), 51–65.
- [106] Olsen, Megan, Harrington, Kyle, and Siegelmann, Hava. Computational emotions in a population dynamics cellular automata encourage collective behavior. In *International Conference on Complex Systems* (June 2011).
- [107] Olsen, Megan, and Siegelmann, Hava. Multi-agent system that attains longevity via death. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)* (January 2007).
- [108] Olsen, Megan, Sitaraman, Ramesh, Siegelmann-Danieli, Nava, and Siegelmann, Hava. Mathematical and computational models for cellular space in cancer growth. In *Proceedings of the American Association for Cancer Research* (April 2010).
- [109] Olsen, Megan M., Siegelmann-Danieli, Nava, and Siegelmann, Hava T. Dynamic computational model suggests that cellular citizenship is fundamental for selective tumor apoptosis. *PLoS One* 5, 5 (May 2010).
- [110] Olsen, M.M., Siegelmann-Danieli, N., and Siegelmann, H.T. Robust artificial life via artificial programmed death. *Artificial Intelligence* 172, 6-7 (2008), 884 – 898.
- [111] Owen, M. R., Alarcon, T., Maini, P. K., and Byrne, H. M. Angiogenesis and vascular remodelling in normal and cancerous tissues. *Journal of Mathematical Biology* 58, 4-5 (2009), 689–721.
- [112] Park, Kihong. The internet as a complex system. In *The Internet as a Large-Scale Complex System*, K. Park and W. Willinger, Eds. Oxford University Press, 2005, pp. 1–89.
- [113] Peirce, Shayn M. Computational and mathematical modeling of angiogenesis. *Microcirculation* 15, 8 (2008), 739–751.

- [114] Peirce, Shayn M., Van Gieson, Eric J., and Skalak, Thomas C. Multicellular simulation predicts microvascular patterning and in silico tissue assembly. *The FASEB Journal* 18, 6 (2004), 731–733.
- [115] Petraglio, E., Henry, J., and Tempesti, G. Arithmetic operations on self-replicating cellular automata. In *Proceedings of the 5th European Conference on Advances in Artificial Life* (1999), D. Floreano, J. Nicoud, and F. Mondada, Eds., pp. 447–456.
- [116] Picard, Rosalind. *Affective Computing*. MIT Press, 1997.
- [117] Plotkin, Joshua, and Nowak, Martin A. Different effects of apoptosis and dna repair on tumorigenesis. *Journal of Theoretical Biology* 214 (2002), 453–467.
- [118] Plutchik, R. The nature of emotions. *American Scientist* 89 (2001), 344–350.
- [119] Prokopenko, Mikhail, Poulton, Geoff, Price, Don, Wang, Peter, Valencia, Phillip, Hoschke, Nigel, Farmer, Tony, Hedley, Mark, Lewis, Chris, and Scott, Andrew. *Self-Organising Impact Sensing Networks in Robust Aerospace Vehicles*. Idea Group, 2006, ch. 7, pp. 189–230.
- [120] Prokopenko, Mikhail, Wang, Peter, Valencia, Philip, Price, Don, Foreman, Mark, and Farmer, Anthony. Self-organizing hierarchies in sensor and communication networks. *Artificial Life* 11, 4 (2005), 407–426.
- [121] Pultack IF, Erff M, Ashkenazi A. Direct stimulation of apoptotic signaling by soluble apo21/tumor necrosis factor-related apoptosis-inducing ligand leads to selective killing of glioma cells. *Clin Cancer Res* 7 (2001), 1362–69.
- [122] Quaranta, V., Weaver, A.M., Cummings, P.T., and Anderson, A.R.A. Mathematical modeling of cancer: The future of prognosis and treatment. *Clinica Chimica Acta* 357 (2005).
- [123] R, Kim, M, Emi, K, Tanabe, Y, Uchida, and K, Arihiro. The role of apoptotic or nonapoptotic cell death in determining cellular response to anticancer treatment. *Eur J Surg Oncol* 32 (2006), 269–77.
- [124] Rolls, E. *What are emotions, Why do We Have Emotions, and What is Their Computational Basis in the Brain?* Oxford University Press, 2005.
- [125] Rosenberg, Arnold L. Cellular automata: Food-finding and maze-threading. In *Proceedings of the 37th international Conference on Parallel Processing. ICPP. IEEE Computer Society, Washington, DC* (2008), pp. 528–535.
- [126] Roth, Fabian, Siegelmann, Hava T., and Douglas, Rodney J. The self-construction and -repair of a foraging organism by explicitly specified development from a single cell. *Artificial Life* 13, 4 (2007), 347–368.
- [127] Russell, James A. A circumplex model of affect. *Journal of personality and social psychology* 39 (1980), 1161–1178.
- [128] Ryoo, Hyung Don, Gorenc, Travis, and Steller, Hermann. Apoptotic cells can induce compensatory cell proliferation through the jnk and the wingless signaling pathways. *Developmental Cell* 7, 4 (2004), 491 – 501.

- [129] S, Sanga, HB, Frieboes, X, Zheng, R, Gatenby, EL, Bearer, and et al. Predictive oncology: A review of multidisciplinary, multiscale in silico modeling linking phenotype, morphology and growth. *Neuroimage* 37 (2007), S120–34.
- [130] Sachs, R.K., Hlatky, L.R., and Hahnfeldt, P. Simple ode models of tumor growth and anti-angiogenic or radiation treatment. *Mathematical and Computer Modelling* 33 (2001), 1297–1305.
- [131] Saidi, Hassen, Dutertre, Bruno, Levy, Joshua, and Valdes, Alfonso. Self-regenerative software components. In *SSRS '03: Proceedings of the 2003 ACM workshop on Survivable and self-regenerative systems* (New York, NY, USA, 2003), ACM Press, pp. 115–120.
- [132] Sanfey, A.G., Rilling, J.K., Aronson, J.A., Nystrom, L.E., and Cohen, J.D. The neural basis of economic decision-making in the ultimatum game. *Science* 300 (2003), 1755–1758.
- [133] Scholz, Christopher H., Sykes, Lynn R., and Aggarwal, Yash P. Earthquake prediction: A physical basis. *Science* 181, 4102 (1973), 803–810.
- [134] Schonfisch, B., and de Roos, A. Synchronous and asynchronous updating in cellular automata. *Biosystems* 51 (1999), 123–143.
- [135] Shalizi, Cosma Rohilla. Methods and techniques of complex systems science: An overview. In *Complex Systems Science in Biomedicine*, Evangelia Micheli-Tzanakou, Thomas S. Deisboeck, and J. Yasha Kresh, Eds., Topics in Biomedical Engineering International Book Series. Springer US, 2006, pp. 33–114.
- [136] Shehory, Onn, Sycara, Katia, Chalasani, Prasad, and Jha, Somesh. Agent cloning: An approach to gent mobility and resource allocation. *IEEE Communications Magazine* (1998).
- [137] Shim, Eun Bo, Kwon, Young-Guen, and Ko, Hyung Jong. Computational analysis of tumor angiogenesis patterns using a two-dimensional model. *Yonsei Medical Journal* 46, 2 (2005), 275–283.
- [138] Shirinifard, Abbas, Gens, J. Scott, and Zaitlen. 3d multi-cell simulation of tumor growth and angiogenesis. *PLoS ONE* 4, 10 (2009), e7190.
- [139] Simpson, Matthew J., Merrifield, Alistair, Landman, Kerry A., and Hughes, Barry D. Simulating invasion with cellular automata: Connecting cell-scale and population-scale properties. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)* 76, 2 (2007), 021918.
- [140] Sipper, M. Co-evolving non-uniform cellular automata to perform computations. *Physica D* 92 (1996), 193–208.
- [141] Sirot, E., and Touzalin, F. Coordination and synchronization of vigilance in groups of prey: The role of collective detection and predators' preference for stragglers. *The American Naturalist* 173 (2009), 47–59.
- [142] SL, Spencer, RA, Gerety, KJ, Pienta, and S, Forrest. Modeling somatic evolution in tumorigenesis. *PLoS Comput Biol* 2 (2006), e108.
- [143] Spinoza, Benedictus de. Ethics part 3.

- [144] Sterritt, Roy, and Hinchey, Mike. *Apoptosis and Self-Destruct: A Contribution to Autonomic Agents?*, vol. 3228 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2004, pp. 262–270.
- [145] Sumpter, David J. T. The principles of collective animal behaviour. *Phil. Trans. R. Soc. B* 361 (2006), 5–22.
- [146] Sumpter, David J.T., and Beekman, Madeleine. From nonlinearity to optimality: pheromone trail foraging by ants. *Animal Behaviour* 66 (2003), 273–280.
- [147] Sutton, Richard S., and Barto, Andrew G. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
- [148] Tilman, David, Lehman, Clarence L., and Kareiva, Peter. *Population Dynamics in Spatial Habitats*. Princeton University Press, NJ, 1997, pp. 3–20.
- [149] Topa, Paweł. Dynamically reorganising vascular networks modelled using cellular automata approach. In *Proceedings of the 8th international conference on Cellular Automata for Research and Industry* (Berlin, Heidelberg, 2008), ACRI '08, Springer-Verlag, pp. 494–499.
- [150] Toyama, K., and Hager, G. D. If at first you don't succeed... In *Proceedings of the 14th National Conference on Artificial Intelligence* (1997).
- [151] Turing, A. M. The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 237, 641 (August 1952), 37–72.
- [152] von Neumann, J. *Theory of Self-reproducing Automata*. University of Illinois Press, Urbana, 1966.
- [153] Wolkenhauer, Olaf, Auffray, Charles, Baltrusch, Simone, Blthgen, Nils, Byrne, Helen, Cascante, Marta, Ciliberto, Andrea, Dale, Trevor, Drasdo, Dirk, Fell, David, Ferrell, James E., Gallahan, Daniel, Gatenby, Robert, Gnther, Ulrich, Harms, Brian D., Herzel, Hanspeter, Junghanss, Christian, Kunz, Manfred, van Leeuwen, Ingeborg, Lenormand, Philippe, Levi, Francis, Linnebacher, Michael, Lowengrub, John, Maini, Philip K., Malik, Arif, Rateitschak, Katja, Sansom, Owen, Schfer, Reinhold, Schrrle, Karsten, Sers, Christine, Schnell, Santiago, Shibata, Darryl, Tyson, John, Vera, Julio, White, Michael, Zhivotovsky, Boris, and Jaster, Robert. Systems biologists seek fuller integration of systems biology approaches in new cancer research programs. *Cancer Research* 70, 1 (2010), 12–13.
- [154] Yamada, Kenneth M., and Cukierman, Edna. Modeling tissue morphogenesis and cancer in 3d. *Cell* 130 (2007).
- [155] Zechmeister, Ingrid, Freiesleben de Blasio, Birgitte, and Garnett, Geoff. Hpv-vaccination for the prevention of cervical cancer in austria: a model based long-term prognosis of cancer epidemiology. *Journal of Public Health* 18 (2010), 3–13.
- [156] Zhang, L., Athale, C.A., and Deisboeck, T. S. Development of a three-dimensional multiscale agent-based tumor model: Simulating gene-protein interaction profiles, cell phenotypes and multicellular patterns in brain cancer. *Journal of Theoretical Biology* (2007).
- [157] Zykov, V., Mytilinaios, E., Adams, B., and Lipson, H. Robotics: Self-reproducing machines. *Nature* 435 (2005), 163–164.