Interdisciplinary Description of Complex Systems 19(4), 493-501, 2021

AN ARTIFICIAL IMMUNE SYSTEM APPROACH TO AUTOMATED PROGRAM VERIFICATION: TOWARDS A THEORY OF UNDECIDABILITY IN BIOLOGICAL COMPUTING

Soumya Banerjee[1, 2, *]

[1]University of Oxford
Oxford, United Kingdom

[2]Ronin Institute
Montclair, USA

DOI: 10.7906/indecs.19.4.3
Regular article

Received: 4 April 2020.
Accepted: 21 December 2021.

ABSTRACT

We propose an immune system inspired Artificial Immune System algorithm for the purposes of automated program verification. It is proposed to use this Artificial Immune System algorithm for a specific automated program verification task: that of predicting shape of program invariants. It is shown that the algorithm correctly predicts program invariant shape for a variety of benchmarked programs. Program invariants encapsulate the computability of a particular program, e.g. whether it performs a particular function correctly and whether it terminates or not. This work also lays the foundation for applying concepts of theoretical incomputability and undecidability to biological systems like the immune system that perform robust computation to eliminate pathogens.

KEY WORDS

artificial immune system, program invariant, undecidability, incomputability, biological computing, immuno-computing, fundamental limits on biological computing

CLASSIFICATION

JEL: I10, O30

*Corresponding author, η: soumya.banerjee@maths.ox.ac.uk; +1 505 277 3122;
Department of Computer Science, 1, University of New Mexico, Albuquerque, NM, 87131, USA

INTRODUCTION

The biological immune system has proved to be a rich source of inspiration for computing [1-15]. Artificial immune systems (AISs) take inspiration from the immune system to provide powerful metaphors for robust and distributed computing. In this article, I employ an immune system inspired approach to solve a problem in program verification: that of finding a program invariant.

An invariant of a program is a mathematical formula that captures the semantics of the program [16] and is used in automatic program verification. The shape of an invariant is its approximate polynomial representation. Once the shape of the invariant is predicted, deterministic techniques can be used to generate the exact form of the invariant [17]. Hence, the prediction of invariant shape is of paramount importance for program verification.

An AIS algorithmic framework is proposed to carry out the machine-learning task of predicting invariant shape from an instance of a program. Program invariants encapsulate the computability of a particular program, e.g. whether it performs a particular function correctly and whether it terminates or not. We hope this work will also lay the foundation for applying concepts of theoretical incomputability and undecidability to biological systems like the immune system that perform robust computation to eliminate pathogens [8-15].

IMMUNOLOGICAL PRELIMINARIES

A chemical species that can be recognized by the adaptive immune system is known as an antigen (*Ag*). When an organism is exposed to an *Ag*, some specialized immune system cells called B cells respond by producing chemicals called antibodies (*Ab*'s). *Ab*'s are molecules attached primarily to the surface of B cells whose aim is to recognize and bind to *Ag*'s. By binding to these *Ab*'s the Ag stimulates the B cell to proliferate and mature into plasma cells that secrete *Ab*. An organism is expected to encounter a given *Ag* repeatedly during its lifetime. The effectiveness of the immune response to secondary encounters is enhanced by the presence of memory cells associated with the first infection, capable of producing high- affinity *Ab*'s after repeat encounters. Such a strategy ensures that the speed and accuracy of the immune response becomes successively higher after each infection. This gives rise to associative memory where the stored pattern is recovered through the presentation of an incomplete version of the pattern. The repertoire of activated B cells is diversified [18-21] and B-cells with higher affinity for the antigen are selected to enter the pool of memory cells.

AUTOMATED PROGRAM VERIFICATION AND PROGRAM INVARIANTS

The field of automated program verification started with seminal work by Floyd [22] and Hoare [23]. They introduced the concept of a loop invariant: a mathematical formula that remains true throughout the execution of a loop. The loop invariant completely captures the semantics of the loop, and along with the program preconditions and postconditions, can be used to show correctness of the program [23].

Previous work [16] has shown how the loop invariant for a particular program can be generated by a priori agreement on the shape of the invariant: the approximate polynomial representation of the invariant. However, the shape of the loop invariant can be hard to deduce for many programs.

The following shows an example program:

```
{A □ 0, B □ 0}
x := A;
y := B;
z := 0;
```

```
while x > 0 do
    if odd(x) then z := z + y;
    y := 2 * y;
    x := x/2;
end while
```

Assuming the shape of the program invariant as $\mathbf{I}_{shape}$: $Ax + By + Cz + Dxy + Eyz + Fxz + Gxyz + H = 0$, (where A, B, C, D, E, F, G and H are constants or program variables), using quantifier elimination [16] the final loop invariant is Ifinal: z + xy –AB = 0. Coupled with a precondition **P**: $\{A \geq 0 \wedge B \geq 0 \wedge x = A \wedge y = B \wedge z = 0\}$ and a postcondition **Q**: $\{z = A \cdot B\}$, it can be shown that this invariant is consistent with **Q** i.e. the program correctly multiplies 2 numbers A and B and stores the result in *z*.

Finding the precise shape of the loop invariant is generally a non-trivial process and the algorithm proposed aims to use 'cues' from the program to make informed predictions about the invariant shape and ultimately help in automated program verification.

PROPOSED COMPUTATIONAL FRAMEWORK

Here we propose a computational framework for predicting program invariants. An AIS algorithm will be used to generate shapes of program invariant. Initially the AIS will be trained on programs, for which the shape of invariant is known. Then a program will be presented to the AIS and it will try to predict the form of the invariant.

An AIS approach presents many advantages over a traditional Machine Learning (ML) approach. In an AIS, recognition can be *sloppy* [24] i.e. if it has previously recognized program *P* (with an invariant *I*), then a new program *P'* 'similar' to *P*, can also be recognized, and an invariant *I'* can be generated (that is similar in form to *I*). This is akin to our immune system recognizing a previously encountered pathogen (program), and generating antibodies (invariant) similar to the previously produced antibodies.

The natural immune system produces antibodies by a process of mutation, and the same process is emulated in AIS algorithms. A candidate solution (invariant) will be generated, and then the solution will be improved by *in-silico* mutation.

Previously encountered programs and their corresponding invariants will be stored as memory B cells. When a program similar to a stored one is presented, the time taken to generate the invariant will be shorter than the time taken to generate the original invariant (*secondary response*).

COMPONENTS OF THE ARTIFICIAL IMMUNE SYSTEM

Here we define the specific components of the AIS have to be determined. What is the program analogue of an antigen and an antibody?

A *program fragment* is defined to be either an assignment statement, a statement containing an iteration construct (for, while, repeat, etc.), or a statement having a conditional check (if <condition> then) e.g. $x := x + 2$, and *while* $(x > 0)$ *do*, and *if* $(x > 3)$ *then*, are all program fragments.

The analogue of an antigen is a program fragment and the corresponding analogue of an antibody is an invariant for the program fragment it recognizes. Hence, the AIS will be presented with an antigen (program fragment), and the immune system cells will either produce the antibody (invariant) immediately if it has encountered this antigen before, or willundergo mutations to generate the correct antibody (invariant).

The individual invariants for each program fragment will then be recombined to generate the invariant for the whole program.

A SHAPE SPACE AND ANTIGENIC DISTANCE FOR PROGRAMS

We need a measure of distance between disparate program fragments, so that the AIS can recognize them and generate an antibody in response. For a natural immune system, the antibody combining region relevant to antigen binding can be specified by a number of 'shape' parameters [25] which denote the size and shape of the combining site or physical characteristics of the amino acids.

If there are *N* shape parameters, they can be combined into a vector, and antibody combining sites and antigenic determinants can be described as points **Ab** and **Ag**, in an *N* - dimensional Euclidean vector-space called *shape space* [25].

Antigenic distance between 2 antigens is the distance in shape space [26] between them e.g. $||\mathbf{Ag_1} - \mathbf{Ag_2}||$ is the distance between antigens $\mathbf{Ag_1}$ and $\mathbf{Ag_2}$ in shape space S. The *antibody distance* is the distance $||\mathbf{Ab_1} - \mathbf{Ab_2}||$ in shape space between 2 antibodies $\mathbf{Ab_1}$ and $\mathbf{Ab_2}$.

I define the *program fragment shape space* as the *N*-dimensional Euclidean vector space of program fragment characteristics like identifier name, exponent on the identifier, operator, etc. I define the corresponding *program fragment antigenic distance* as the distance $||\mathbf{P_1} - \mathbf{P_2}||$ between 2 program fragments $\mathbf{P_1}$ and $\mathbf{P_2}$ in program fragment shape space. The *program fragment antibody (invariant) distance* is the distance $||\mathbf{I_1} - \mathbf{I_2}||$ between 2 program fragments $\mathbf{I_1}$ and $\mathbf{I_2}$ in program fragment shape space.

Let us consider 2 program fragments $\mathbf{P_1}$: *x := x + 2* and $\mathbf{P_2}$: *t := t + 2*. The corresponding antibody (invariant) for $\mathbf{P_1}$ is $\mathbf{I_1}$: *x = x + 2n*, where *n* is a program variable or constant (since upon *n –1* iterations, *x* gets the value *x + 2n*). Let $\mathbf{P_1}$ and $\mathbf{I_1}$ constitute the training set. Then the AIS should be able to produce an antibody (invariant) for the program fragment $\mathbf{P_2}$ even though it has never encountered this antigen (program) before. The correct invariant is $\mathbf{I_2}$: *t = t + 2n* (where *n* is a program variable or constant) and this is indeed what the AIS generates by somatic hypermutation. The program $\mathbf{P_1}$ differs from $\mathbf{P_2}$ by 1 mutation (replacing *x* by *t* onboth sides of the assignment) i.e. the *program fragment antigenic distance* $||\mathbf{P_1} - \mathbf{P_2}||$ is 1. The invariants $\mathbf{I_1}$ and $\mathbf{I_2}$ also differ by 1 mutation (replacing *x* by *t*) i.e. the *program fragment antibody (invariant) distance* $||\mathbf{I_1} - \mathbf{I_2}||$ is 1. Hence, when an AIS trained on ($\mathbf{P_1}$, $\mathbf{I_1}$) ispresented with $\mathbf{P_2}$, it produces $\mathbf{I_2}$ using one mutation from $\mathbf{I_1}$ (Fig. 1).

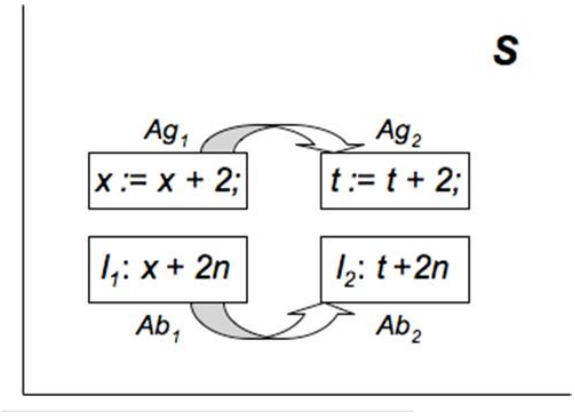

Figure. 1. AIS mutation from the assignment statement $\mathbf{Ag_1}$ (x := x + 2;) and invariant $\mathbf{Ab_1}$ (x + 2n) to $\mathbf{Ag_2}$ (t := t + 2;) and invariant $\mathbf{Ab_2}$ (t + 2n) in shape space S.

PROPOSED ALGORITHM

In this section we outline the proposed immune system inspired algorithm. The AIS would be trained on the antigen (program fragment) $\mathbf{P_1}$: *x := x + 2*and given the antibody (invariant)

$\mathbf{I_1}$: $x = x + 2n$ as a solution (*training phase*). The AIS stores the solution $\mathbf{I_1}$ as a memory detector.

When an entire program (as opposed to a program fragment) is presented to the AIS, it breaks the program up into program fragments (all the assignment statements in the program), and then 'presents' each of these antigens (fragments) to itself.

If an antigen (program fragment) $\mathbf{P_2}$ 'similar' to $\mathbf{P_1}$ is detected, it will generate $\mathbf{I_1}$ as a candidate solution. If $\mathbf{I_1}$ itself does not act as an invariant, the AIS will keep on carrying out randomly on $\mathbf{I_1}$ until it evolves the final antibody (invariant) $\mathbf{I_2}$ that will act as the invariant for the program presented (*somatic hypermutation phase*). This is akin to how the natural immune system mutates B cell receptors and ultimately produces a receptor that can recognize the antigen. The algorithm may also use some heuristics to guide the mutation process e.g. if an antigen (program fragment) of the form $p := p + 5$ is encountered, it would search its repertoire for a program fragment that is closest in program shape space to this e.g. $x := x + 5$ is closer to the presented antigen (1 mutation) than $y := y + 7$ (2 mutations). Additionally, we will have to ensure that each mutation is *sound* i.e. there is no such mutation that would generate a wrong invariant for the corresponding mutated program fragment. In the last step, the AIS incorporates $\mathbf{I_i}$ into its memory pool (*learning phase*).

The AIS then presents the next program fragment $\mathbf{P_3}$, generates the invariant $\mathbf{I_3}$ and stores it in the memory population, and so on until all program fragments have been presented. Finally, the AIS combines all invariants linearly, producing a polynomial (shape of invariant)that captures the semantics of the entire program.

RESULTS

The AIS (trained on $\mathbf{P_1}$, $\mathbf{I_1}$) presented with suites of entire programs would successfully generate the shape of the invariant. The first program is shown below:

```
(x,y,u,v) := (a,b,b,0);
x := a; y := b;
u := b; v := 0;

while (x □ y) do
    while (x > y) do x := x - y; v := v + u; end while;
    while (x < y) do y := y - x; u := u + v; end while;
end while
```

This program takes 2 positive integers *a* and *b*, and calculates their greatest common divisor and least common multiple. The AIS presents itself with each assignment statement sequentially. The first 4 assignment statements (lines 1-2) have no invariant, since they are not contained inside any loop. Hence, the AIS does not generate any invariant for them. The progress of the algorithm on the next 2 assignment statements (x := x –y; v := v + u), Fig. 2.

The AIS starts from the training set (P1: x := x + 2 & I1: x = x + 2n) and then mutates the operators and operands to create the invariant I3: x = x –yn for the program fragment P3: x := x –y. The AIS stores I3 in the memory population and for the next assignment statement (v := v + u;), it starts mutating from (P3, I3) until it creates the invariant I4: v = v + un for the program fragment P4: v := v + u.

For the next set of assignment statements (y := y –x; u := u + v;), the AIS then generates the invariants $\mathbf{I_5}$: $y = y - xn$ and $\mathbf{I_6}$: $u = u + vn$ (not shown). The 4 invariants $\mathbf{I_1}$, $\mathbf{I_2}$, $\mathbf{I_3}$ & $\mathbf{I_4}$ are then combined linearly (with *n* being substituted for all program variables, namely *x, y, u, v*) to

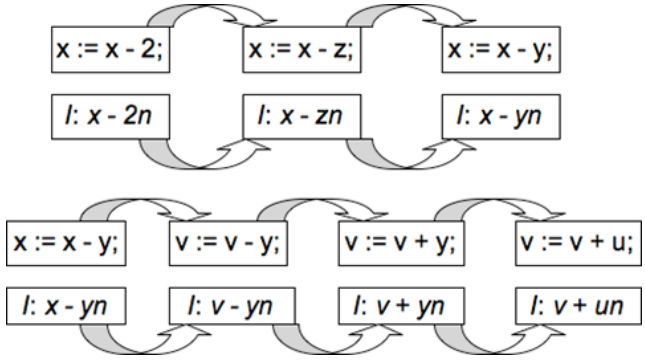

Figure. 2. AIS mutations for the assignment statements x := x –y; v := v + u.

yield the invariant shape $\mathbf{I_{shape}}$: $Ax + Bv + Cy + Du + Exy + Fy^2 + Guy + Hvy + Jxu + Ku^2 + Lvu + Mx^2 + Nvx + Pv^2 + Q = 0$, where $A, B, C, D, E, F, G, H, J, K, L, M, N, P$ and Q are constants or program variables. This is the correct invariant shape, since using quantifier elimination [16], the final invariant yielded is $\mathbf{I_{final}}$: $xu + yv - ab = 0$ (with $A = B = C = D = E = F = G = K = L = M = N = P = 0, Q = -ab, H = J = 1$).

Finally we test the AIS on another standard program [16] shown below:

```
{A ≥ 0, B ≥ 0}
x := A;
y := B;
z := 1;

while y > 0 do
    if odd(y) then y := y - 1; z := x * z;
    else x := x * x;    y := y/2;
end while
```

This program calculates A^B and stores it in z. The AIS would calculate the invariant for the program fragment $\mathbf{P_5}$: $z := x * z$ as $\mathbf{I_5}$: $z = x^n * z$. For the program fragment $\mathbf{P_6}$: x := x * x, it generates the invariant $\mathbf{I_6}$: $x = exp(x, exp(2,n))$, where $exp()$ is the exponentiation function. Combining all the program fragment invariants, gives us the following invariant shape:
$\mathbf{I_{shape}}$: $Azx^x + Bzx^y + Czx^z + D.exp(x,exp(2,x)) + E.exp(x,exp(2,y)) + F.exp(x,exp(2,z)) + G = 0$.

This is the exact shape of the invariants, since quantifier elimination yields the final invariant $\mathbf{I_{final}}$: $zx^y = A^B$ (with $A = C = D = E = F = 0, G = -A^B$).

We can now readily verify the working of the program. When the loop terminates, the invariant is true and $y = 0$, which yields the correct postcondition: $z = A^B$.

The proposed algorithm would use a sequence of mutations, guided by heuristics, to generate the correct invariant for a program invariant.

CONCLUSION AND FUTURE WORK

We have proposed a computational framework for an immune system inspired approach for automated program verification. The immune system inspired algorithm breaks up a program into fragments and presents them to itself. It then generates an invariant in response to each program fragment and ultimately combines them to create the general shape of the invariant. We show how this approach can be used to generate the general form of the program invariant for non-trivial benchmark programs [16].

Future work will focus on theoretical research into whether there are classes of programs

for which a linear combination of individual program fragment invariants might not generate the invariant for the entire program. Another avenue of future investigation would be to look into how mutations on exponentiation would affect the invariant e.g. $x := x + 2$ getting mutated to $x := x^2 + 2$. Lastly, our approach does not consider program fragments having iteration constructs like *while*, *repeat*, etc. and future research will investigate how incorporation of such program fragments can enhance the predictive power of the algorithm.

A lot of work has been done on incomputability, undecidability and program termination in theoretical computer science. The best characterization of this comes in the form of the Halting Problem formulated by Alan Turing. Biological systems also perform computing, e.g. the immune system computes the most efficient way to eliminate pathogens in a timely manner without harming the host [8-15]. However it has been more difficult to define incomputability and undecidability for biological systems.

Program invariants encapsulate the computability and correctness of a particular program, e.g. what it does and whether it terminates or not. This work lays the foundation of applying computability to biological systems especially the immune system that performs computation.

The present work also applies immune system inspired algorithms to find program invariants and prove correctness and termination. It is intriguing to speculate that it is also possible to go in the reverse direction and translate the complexities of the immune system into an equivalent computer program. The translated computer program can then be analyzed for mathematical properties of what it computes [27, 28]. Hence this work can be extended to provide a theoretical framework for understanding the limits of computation in the immune system.

The present computational framework can be used to account for cases when the immune system fails to clear infections as is the case in certain virulent infections [29].

This approach can also be similarly extended to analyse substrates for computing that are non-silicon based and can be used to probe the computational nature of life itself [15].

In summary, the present work applies the theoretical concepts of undecidability to immuno- computing and possibly biological computing in general. We view this work as the first step towards elucidating the fundamental limits of computing in immunology and possibly biology as well.

ACKNOWLEDGEMENTS

The author wishes to thank Dr. Sara Jane-Dunn, Dr. Boyan Yordanov, Prof. Deepak Kapur and Dr. ThanhVu Nguyen for helpful comments.

REFERENCES

[1] de Castro, L.N. and Timmis, J.: *Artificial Immune Systems: A New Computational Intelligence Approach.*
Springer-Verlag, London, 1996,

[2] Hunt, J.E. and Cooke, D.E.: *Learning using an artificial immune system.*
Journal of Network and Computer Applications **19**(2), 189-212, 1996,
http://dx.doi.org/10.1006/jnca.1996.0014,

[3] Dasgupta, D.: *Artificial Immune Systems and Their Applications.*
Springer-Verlag, Berlin, 1999,

[4] Hofmeyr, S.A. and Forrest, S.: *Immunity by design: An artificial immune system.*
Proceedings Genetic and Evolutionary Computation conference, pp.1289-1296, 1999,

[5] de Castro, L.N. and Von Zuben, F.J.: *Artificial Immune Systems: Part I. Basic Theory and Applications.*
School of Electrical and Computer Engineering, University of Campinas, Campinas, 1999,

[6] de Castro, L.N. and Von Zuben, F.J.: *Artificial Immune Systems: Part II – A Survey of Applications.*
School of Electrical and Computer Engineering, University of Campinas, Campinas, 2000,
[7] de Castro, L.N. and von Zuben, F.J.: *Learning and Optimization Using the Clonal Selection Principle.*
IEEE Transactions on Evolutionary Computation **6**(3), 239-251, 2002,
http://dx.doi.org/10.1109/TEVC.2002.1011539,
[8] Banerjee, S. and Moses, M.: *Scale Invariance of Immune System Response Rates and Times: Perspectives on Immune System Architecture and Implications for Artificial Immune Systems.*
Swarm Intelligence **4**, 301-318, 2010,
http://dx.doi.org/10.1007/s11721-010-0048-2,
[9] Banerjee, S.: *Scaling in the immune system.* Ph.D. Thesis.
University of New Mexico, Albuquerque, 2013,
[10] Banerjee, S., et. al.: *The Value of Inflammatory Signals in Adaptive Immune Responses.*
The 10th International Conference on Artificial Immune Systems, 2011,
[11] Drew, L., et al.: *A spatial model of the efficiency of T cell search in the influenza-infected lung.*
Journal of Theoretical Biology **398**(7), 52-63, 2016,
http://dx.doi.org/10.1016/j.jtbi.2016.02.022,
[12] Banerjee, S.: *A Biologically Inspired Model of Distributed Online Communication Supporting Efficient Search and Diffusion of Innovation.*
Interdisciplinary Description of Complex Systems **14**(1), 10-22, 2016,
http://dx.doi.org/10.7906/indecs.14.1.2,
[13] Banerjee, S.; van Hentenryck, P. and Cebrian, M.: *Competitive dynamics between criminals and law enforcement explains the super-linear scaling of crime in cities.*
Palgrave Communications **1**, No. 15022, 2015,
http://dx.doi.org/10.1057/palcomms.2015.22,
[14] Banerjee, S. and Hecker, J.P.: *A Multi-Agent System Approach to Load-Balancing and Resource Allocation for Distributed Computing.*
Complex Systems Digital Campus, World e-Conference, Conference on Complex Systems, 2015,
[15] Banerjee, S.: *A Roadmap for a Computational Theory of the Value of Information in Origin of Life Questions.*
Interdisciplinary Description of Complex Systems **14**(3), 314-321, 2016,
http://dx.doi.org/10.7906/indecs.14.3.4,
[16] Kapur, D.: *Automatically Generating Loop Invariants Using Quantifier Elimination.*
IMACS International IMACS Conference on Applications of Computer Algebra, 2004,
[17] Rodriguez, E. and Kapur, D.: *Automatic Generation of Polynomial Loop Invariants: Algebraic Foundations.*
International Conference on Symbolic and Algebraic Computation, 2004,
[18] Berek, C. and Ziegner, M.: *The maturation of the immune response.*
Immunology Today **14**(8), 400-404, 1993,
http://dx.doi.org/10.1016/0167-5699(93)90143-9,
[19] George, A.J.T. and Gray, D.: *Receptor editing during affinity maturation.*
Immunology Today **20**(4), 196, 1999,
http://dx.doi.org/10.1016/S0167-5699(98)01408-X,
[20] Nussenzweig, M.C.: *Immune receptor editing: Revise and select.*
Cell **95**(7), 875-878, 1998,
http://dx.doi.org/10.1016/S0092-8674(00)81711-0,
[21] Tonegawa, S.: *Somatic generation of antibody diversity.*
Nature **302**(14), 575-581, 1983,
http://dx.doi.org/10.1038/302575a0,

[22] Floyd, R.W.: *Assigning meanings to programs.*
Proceedings of the American Mathematical Society Symposia on Applied Mathematics **19**, 19-31, 1967,

[23] Hoare, C.A.R.: *An axiomatic basis for computer programming.*
Communications of the ACM **12**(10), 576-585, 1969,
http://dx.doi.org/10.1145/363235.363259,

[24] Perelson, A.S. and Wiegel, F.W.: *Some design principles for immune system recognition.*
Complexity, 1999,

[25] Perelson, A.S. and Oster, G.F.: *Theoretical Studies of Clonal Selection: Minimal Antibody Repertoire Size and Reliability of Self-Non-self Discrimination.*
Journal of Theoretical Biology **81**(4), 645-67, 1979,

[26] Smith, D.J.; Forrest, S.; Hightower, R.R. and Perelson, S.A.: *Deriving shape space parameters from immunological data.*
Journal of Theoretical Biology **189**(2), 141-150, 1997,
http://dx.doi.org/10.1006/jtbi.1997.0495,

[27] Dunn, S.J., et al.: *Defining an essential transcription factor program for naïve pluripotency.*
Science **344**(6188), 1156-1160, 2014,
http://dx.doi.org/10.1126/science.1248882,

[28] Yordanov, B., et al.: *A method to identify and analyze biological programs through automated reasoning.*
npj Systems Biology and Applications **2**, No. 16010, 2016,
http://dx.doi.org/10.1038/npjsba.2016.10,

[29] Banerjee, S. et al.: *Estimating Biologically Relevant Parameters under Uncertainty for Experimental Within-Host Murine West Nile Virus Infection.*
Journal of the Royal Society Interface **13**(117), No. 20160130, 2016,
http://dx.doi.org/10.1098/rsif.2016.0130.