History of RSB Interview: Hanoch Gutfreund

December 14, 2020, 8:30-10:00am (EST). Final revision: February 8, 2021

Interviewers:
Patrick Charbonneau, Duke University, patrick.charbonneau@duke.edu
Francesco Zamponi, ENS-Paris

Location:
Over Zoom, from a conference room at the Hebrew University of Jerusalem, Israel.

How to cite:
P. Charbonneau, *History of RSB Interview: Hanoch Gutfreund*, transcript of an oral history conducted 2020 by Patrick Charbonneau and Francesco Zamponi, History of RSB Project, CAPHÉS, École normale supérieure, Paris, 2021, 16 p.
https://doi.org/10.34847/nkl.1adb9r42

PC: Professor Gutfreund, thank you very much for agreeing to sit down with us. As we discussed, the goal of this interview is to talk mostly about the period 1975 to 1995. But to get to understand this period, I'd like to go back a little in your career. To situate us, please help us understand what drew you to physics. How did you get to choose to do a PhD in physics in the first place?

HG: What drew me into physics... I was a student at the Hebrew University. This is the Racah Institute of Physics[1], where [was] the legendary Giulio Racah[2], who was forced to leave Mussolini's Italy because he was a Jew. He settled in Jerusalem in 1940, and he established theoretical physics in Israel. He was a pioneer in applying group theoretical methods to quantum physics and, specifically, to atomic spectroscopy. I was fortunate, when I enrolled as a student at the university in 1957, to have all my theoretical courses from him. Then, it was natural to continue my studies to higher degrees. I did my PhD on nuclear reactions at the Hebrew University[3], and then I went to Stanford as a postdoctoral student of Felix Bloch[4]. That brought me into the field of superconductivity. After a year, Felix Bloch

[1] Racah Institute of Physics : https://en.wikipedia.org/wiki/The_Racah_Institute_of_Physics
[2] Giulio Racah: https://en.wikipedia.org/wiki/Giulio_Racah
[3] H. Gutfreund, מכניזם של ריאקצית הלכידה הקרינתית של נוקלאונים באנרגיות בינוניות (Mechanisms of the radiative capture reaction for intermediate energy nucleons), PhD Thesis, Thesis, Hebrew University of Jerusalem (1966). https://huji-primo.hosted.exlibrisgroup.com/permalink/f/att40d/972HUJI_ALMA21145782460003701
[4] Felix Bloch: https://en.wikipedia.org/wiki/Felix_Bloch

went on a sabbatical. I then began to work with William (Bill) Little[5]. I don't know if you know the name. He is best known for his ideas, or his quest for high temperature superconductivity, specifically organic superconductivity, long before it was discovered, though in a different context. I did a lot of work with him. In 1973-1974, I came back to Stanford for one year, my sabbatical year. At that time Little was also interested in neural networks.

PC: Before we dive into this, do you mind stepping back one second? As you said, your PhD was in nuclear physics and then you went in a clearly solid state physics group. What was your motivation?

HG: To solid state physics?

PC: Yes. What drew you to that?

HG: The main motivation was that Felix Bloch agreed to take me as a postdoctoral student. That was enough. Felix Bloch, anything he would have done and suggested to me, I would [have] followed.

PC: So you knew of his work…

HG: I was close to him and inspired by him.

PC: Getting back to your sabbatical at Stanford…

HG: [0:03:32] I came back to Jerusalem, 1968-69, a young lecturer. I came back to Stanford during the summers and then, in 1973-4, for a sabbatical year to continue my work with Bill Little. We did a lot of work together on the prospect of organic superconductivity and organic conductivity. At that time, Bill Little had an idea of a kind of neural network model[6].

We mark the genesis of neural networks as a field of study with the work of McCulloch and Pitts, in 1943[7], when they introduced a model of binary neurons, connected by so-called synaptic junctions. This was a kind of caricature neuron, but they could show that a network of such binary elements that is connected by a kind of efficacy connections could perform

---

[5] William A. Little (1929-). Daniel Hartwig, "Guide to the William A. Little Papers," Stanford University. Libraries. Department of Special Collections and University Archives (2013). https://oac.cdlib.org/findaid/ark:/13030/c83x88bk/

[6] William A. Little, "The existence of persistent states in the brain," *Math. Biosci.* **19**, 101-120 (1974). https://doi.org/10.1016/0025-5564(74)90031-5

[7] Warren S. McCulloch and Walter Pitts. "A logical calculus of the ideas immanent in nervous activity." *The bulletin of mathematical biophysics* **5,** 115-133 (1943). https://doi.org/10.1007/BF02478259

any calculation that a Turing machine could perform. This was a first implementation of Turing's idea of thinking machines. There were several developments in this direction. There was the idea of a perceptron based on a layered structure of such networks, where you can design calculating devices that would give you a desired connection between input and output. The second major development in this area was the work of psychologist Donald Hebb in 1949[8], who conjectured that every cognitive event that happens in our brain is connected to some persistent activity of a certain class, a certain group of neurons. He also conjectured—and that is a basic paradigm until today—that learning proceeds by changes in the connections between neurons, by changes in synaptic efficacies. This were the two basic elements in the paradigm which developed around the Hopfield model. But, before that one should mention the work of Minsky-Papert on perceptrons[9], which are such layered networks, where a calculation proceeds from input layer to an output layer. Then, there were the first steps to try to model these networks of binary neurons as physical systems of binary, two-state, elements. The classical candidate for that was a system of magnetic dipoles, spins. There was an attempt, already in 1954, to do something like that. Cragg and Temperley are the names associated with this development[10].

In 1974, Bill Little had an idea, which really struck me and had an impact on my thinking for years to come. He suggested that states, namely persistent activities of the group of neurons that would represent a cognitive event in memory—an idea, a command to your muscles—that occurs in the brain could be associated with, could be modeled by a persistent state of a spin system. His idea was the following. He proposed a model of a two-dimensional Ising spin system, constructed of successive layers. In the first layer, all the spins—or all the neurons—represent the activity of those neurons at time t. The second layer represents the activity of the same group of neurons at time t+1, and so on, from layer to layer. Between the layers, you have all the J_{ij} connections from one layer to another. The idea was to explore under what circumstances, as you proceed in time, a persistent activity would develop, *i.e.*, you would have the same configuration from one state to another. The basic element that he introduced was the transition matrix. The transition matrix is the matrix that tells you what is the probability that having a configuration j in one layer would take you to a configuration i in another layer. This, of course, is a huge matrix. Little's

[8] Donald O. Hebb, *The Organization of Behavior: A Neuropsychological Theory* (New York: John Wiley and Sons, 1949).

[9] M. Minsky and S. Papert, *Perceptrons: an introduction to computational geometry*, (Cambridge: MIT Press, 1969). https://en.wikipedia.org/wiki/Perceptrons_(book)

[10] B. G. Cragg and H. N. V. Temperley, "The organisation of neurones: a co-operative analogy" *Electroencephalogr. Clin. Neurophysiol.* **6**, 85-92 (1954). https://doi.org/10.1016/0013-4694(54)90008-5

model introduced for the first time, the notion of temperature into the neural activity. Temperature is an important element in a physical system. The temperature in the neural network was related to the uncertainty of a neuron firing once it received the input from all the other neurons connected to that neuron. He was the first who introduced the notion of temperature into the network. (I shall tell you in a minute what was still missing in Little's model and was introduced later by Hopfield.) Then he showed that if this transition matrix would have degenerate eigenvalues—value 1 separated from the all the other eigenvalues—then a persistent state would develop. He tried to push it a long way.

When I came back to Jerusalem…

PC: Before we do this, and because you knew Bill Little well, do you know what got him interested in neural networks? From superconductivity it seems like a big jump, but maybe it's obvious.

HG: [0:11:52] It's hard to tell, but I can tell you one thing. Having known him well—he's still alive, by the way, at Stanford—he was very imaginative. Whenever we met, whenever we talked, he had some idea that I couldn't tell where it came from. He had a kind of neural network proof of one of Fermat's theorem, for example[11]. Just something very surprising, unexpected. He was imaginative. I don't know what really triggered him. So he invested a lot of effort in his attempts to push his neural networks model.

I came back to Jerusalem. I was fascinated by this. I had a colleague in Jerusalem. His name was Daniel Amit[12]. (Daniel Amit died in 2007.) Together we explored this idea. Little came to Jerusalem in 1980 and we talked about his model, but we did not know how to connect these persistent states, or the configuration in a persistent state, with the eigenvectors and eigenvalues of that transition matrix, which were the indicators if such states existed or not. Then, in 1982 something happened, which for us was an eye opener. It was a revelation. That was John Hopfield's model in 1982[13]. You see, Bill Little went half of the way. He introduced temperature into models of neural networks. It was not there before. Not in the perceptron, neither in the model of McCullogh and Pitts. Hopfield introduced energy. Energy was the parameter that drove the dynamics of such a network.

[11] H. Gutfreund and W. A. Little, "Physicist's proof of Fermat's theorem of primes," *Am. J. Phys.* **50**, 219 (1982); https://doi.org/10.1119/1.12858

[12] Damiel Amit: https://en.wikipedia.org/wiki/Daniel_Amit

[13] John J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proc. Natl Acad. Sci. U.S.A* **79**, 2554-2558 (1982). https://doi.ord/10.1073/pnas.79.8.2554

Let me tell you a little bit about Hopfield's paradigm, summarizing briefly the main ideas, and then continue with my personal and institutional history. My institution, The Hebrew University in Jerusalem is closely related to what happened in those years. We closely collaborated with colleagues at École normale supérieure and a few other places.

Until Hopfield, discussions of neural networks were based on a layered architecture. The calculation of performance of such a network was simple. It proceeded from one layer to the next layer, until the final layer, where you read out the result. Hopfield assumed a network with a lot of feedback. The extreme case would be that every neuron is connected to every other neuron, and then the whole system evolves under a certain dynamical rule. He assumed that the connections between the neurons—what in a neural network are the synaptic efficacies—are symmetrical. This was a great advantage. It was an advantage because it allowed to apply ideas and methods from physics, to develop models, to solve models. It was also a great step backwards, because it took us further away from the real system. In a real neural network, in a living organism, the connections between neurons are certainly not symmetrical. More than that, if neuron *i* effects neuron *j* by a synaptic connection, there is no reason—and usually it's not the case—that neuron *j* effects also neuron *i*. Real networks are really asymmetrical. This was a cause for many debates and also for a lot of criticism from our biological colleagues. They would ridicule us for making assumptions which are so clearly remote from any real system. Still this symmetry, which one could relax afterwards—I'll tell you later how that happened—played an essential role in the early days of the Hopfield model. Hopfield, already in his seminal paper, in the last sentence, says that these systems are similar to spin glass systems, not more than that. But Daniel Amit and myself, we immediately knew how and in what direction to proceed.

At that time, a bright Israeli student came back to Israel after his post-doctoral work with Bert Halperin in Harvard[14]. His name is Haim Sompolinsky[15]. (Haim Sompolinsky, until today, is a star and I would say a superstar of this field.) He came back and he joined another university in Israel, Bar-Ilan University. He knew a lot about spin glasses at that time, more than any one of us. We invited him to Jerusalem. We started to discuss these things. He was immediately taken by it. The three of us, over the next few years had the greatest time. For me it was the most exciting period in my life as a scientist. This was a time when there was a sense of pioneering something new and profound. There was a real sense of great challenges, and one result followed another.

[14] Bertrand I. Halperin: https://en.wikipedia.org/wiki/Bertrand_Halperin
[15] Haim Sompolinsky: https://en.wikipedia.org/wiki/Haim_Sompolinsky

The idea was the following. You have a network of such binary elements that are connected by the synaptic efficacies. These connections are random and since initially every element is connected to all other elements we had a system with long-range interactions. In a way, there was already a model—which was a landmark in spin glass research, which was very intensive in those days—and that was the Sherrington-Kirkpatrick model. This was a spin glass with random connections, with long-range interactions. The long-range interactions guaranteed that one could use mean-field theory to treat such systems, and, of course, the randomness of connections implied that in order to do averaging over this disorder you have to use what is called the replica trick. The replica trick at that time was already developed, discussed, a main topic in the study of disordered systems, by people like Giorgio Parisi in Rome and others. The elements were already there. The main idea of this model, applied to neural networks, was that such a network is exposed to an external stimulus which would impose a certain initial configuration. That configuration would evolve, and the question is: whether under the specific dynamics, that configuration would, after a certain amount of steps, end up in a certain persistent configuration. Those persistent configurations were called attractors, and this whole scheme was called "calculating by attractors". The goal was to find a system in which one could embed a large amount of such persistent states, namely configurations, such that in the end of a certain dynamical process the system would converge and then oscillate (because of the uncertainty introduced by the temperature) in the very close neighborhood of one of such states. The questions were: How many such states could you embed in a system, namely, what is its storage capacity? What would be their basins of attraction? This was the essence of the model.

We were able—Daniel Amit, Haim Sompolinsky and myself—to find an analytical solution of the Hopfield model, where we assumed that certain configurations, which were chosen in advance, would then serve as the attractors[16]. They represented the cognitive events of the neural network. Initially, those configurations were chosen at random, and without any correlation between them. They were also unbiased, namely, the same number of plus and minus directions of the representative spins in that configuration. Then, the idea was to find the symmetric efficacies that connect every two elements, the neurons, namely the J_{ij} parameters, that would then drive the dynamics, when the system was exposed to any external stimulus, namely an initial configuration, to flow to one of those attractors,

[16] Daniel J. Amit, Hanoch Gutfreund and Haim Sompolinsky, "Storing infinite numbers of patterns in a spin-glass model of neural networks," *Phys. Rev. Lett.* **55**, 1530 (1985). :https://doi.org/10.1103/PhysRevLett.55.1530

to one of those configurations. That was the challenge. We called our success to solve this challenge a great triumph of statistical physics, because that was the first time that statistical physics was introduced brain research (You may ask a question because I have a lot more to say.)

PC: Were you at any point in this process engaged in a conversation with Hopfield, or were you really doing this on your own?

HG: [0:27:18] Yes, Hopfield knew all about that. I shall tell you in a minute about the international connections, collaborations, the major meetings, the other players at the time when a lot of that happened.

We published our work in 1985. In 1986, the Institute for Theoretical Physics in Santa Barbara[17], Robert Schrieffer[18]—one of the BCS trio in superconductivity—was the director, and he then organized a special program on neural networks. John Hopfield was the head of that program. We all participated there. The interactions between the major players in the field were very intensive in those days. There was a very strong group in École normale supérieure, with people like Gérard Toulouse[19], Marc Mézard and few others. There was Giorgio Parisi, Miguel Virasoro, Luca Peliti in Rome. And there were these other groups that continued to do the layered neural networks studies with artificial intelligence applications. These were, at the same time, the first days of AI. I remember the activity in the United States around an effort to apply these layered neural networks to the mechanical recognition of zip codes[20]. There was a whole drama associated with that. We were not involved in that effort but we followed it with interest and curiosity. When we came back from Santa Barbara, we organized at the Institute for Advanced Studies in the Hebrew University[21], in 1987, a group like the one in Santa Barbara with scholars from physics, neural networks, computation, and biology, we had representatives of all four disciplines. I will tell you something else about what happened during that year, in '87, in Jerusalem, because it's another landmark in the evolution of this field. Let me come back to this a little later.

---

[17] John Hopfield and Peter Young, "Spin Glasses, Computation, and Neural Networks" September to December 1986 Institute for Theoretical Physics, University of California at Santa Barbara. See, e.g., Dana H. Ballard. "Modular learning in neural networks" In: *Proceedings of the sixth National conference on Artificial intelligence – Vol. 1* (AAAI'87). AAAI Press, 279–284 (1987).
[18] John Robert Schrieffer was director of the Institute of Theoretical Physics from 1984 to 1989: https://en.wikipedia.org/wiki/John_Robert_Schrieffer
[19] Gérard Toulouse: https://en.wikipedia.org/wiki/G%C3%A9rard_Toulouse
[20] See, *e.g.*, LeNet: https://en.wikipedia.org/wiki/LeNet
[21] Israel Institute of Advanced Studies: https://en.wikipedia.org/wiki/Israel_Institute_for_Advanced_Studies

Let me now tell you more about the Hopfield model. The Hopfield model had several advantages. The first is that it could be solved analytically; once you have a model that can be solved it gives you a certain confidence in what you are doing. Then, that model was amenable to numerous modifications, relaxing all kinds of the initial restrictions. Our main result was that we could compute the storage capacity of such a model. Then, we could relax many of the initial limitations. For example, we could show that our results are very robust with respect to all kinds of changes. Introducing asymmetry in the connective junctions just acts as an additional noise—similar to temperature—another source of noise. It can be perceived as an effective temperature. In fact, we could show that temperature noise, including noise introduced by asymmetry, has certain beneficial effects. Because in the original Hopfield model, in addition to those basins of attraction which represent those final states that presumably represent certain cognitive events, there are many spurious states which have nothing to do with the "meaningful" attractors. Noise—including noise due to asymmetry—erases those spurious attractors, makes the energy surface more smooth, leaving the deep meaningful attractors, eliminating the others, particularly the exponentially many spin glass minima, which are still there in the spin system. So a lot could be done.

I still did not mention replica symmetry, or the replica method. Clearly, when you deal with such disordered systems, with quenched disorder and you want to calculate the free energy you have to average over that disorder. The free energy is expressed by the logarithm of the partition function function *Z*. It's not simple to average over the logarithm of the partition function, and therefore one applies this replica method, the replica trick, which is mathematically very elegant and mathematically raises deep questions, some of them are still unresolved, but the method works. Very amply supported in cases where you can reach the results by other methods, very amply supported by numerical calculations. Haim Sompolinsky was familiar with the replica method at the time when he joined us, but for us it was also a learning period, a very rewarding learning period, coping with all these mathematical tools. Haim Sompolinsky, as I told you, came back to Bar-Ilan University. Our dream and our desire was that he joins us at the Hebrew University. That was our wish. That happened very fast, it was also his wish and Haim Sompolinsky has been at the university ever since, and is still very active there.

FZ: If I understand well, it was Sompolinsky who brought the replica method…

HG: Personally to us, yes.

FZ: Was Daniel Amit not already aware of the method? I think he was already in contact with Giorgio Parisi at the time…

HG: [0:36:29] Yes. Daniel Amit was in contact with Giorgio Parisi. The last few years of his life he spent in Rome with Parisi. There is this Daniel Amit memorial book of which Parisi is one of the editors[22]. He beautifully described there their previous interactions. Daniel before that did not work on spin glasses. I certainly did not. I did other things. For us, it was a very rewarding learning period.

The next thing that happened—and that was a monumental breakthrough—was in 1987. I have already mentioned that we organized a group on Neural Networks and the emerging field of computational neuroscience at the Institute for Advanced Studies in Jerusalem. Many of the players in the field participated were here either for the whole year or for short periods. Gérard Toulouse from École normal supérieure was here for the whole year. At that time, there were two possible directions for neural networks research to follow. One was artificial intelligence, which means the design and study artificial device systems of such neural networks to mimic systems of memory, of learning, of recognition, of generalization, of categorization. The other one was to try to make better contact with biology.

The choice between these two options was debated, and ultimately the main effort in Jerusalem evolved in the direction of biology. I don't know if it was a formal decision. We were fortunate that one of the scholars at the medical school did brain research and was attracted to research on neural networks. His name is Moshe Abeles[23]. He was one of the pioneers to recognize the importance of theoretical biology. In those days, we had bitter debates, controversies, with our biological and medical colleagues at our university. They thought that it was a waste of time and money, that the only way that biology, brain research and research on biological networks could proceed was through experimental studies of specific systems and functions. They were very suspicious of our theoretical approach to that. Moshe Abeles was maybe the only one who strongly believed already then that biology would make progress, like other exact sciences, by recourse to theoretical methods.

Moshe Abeles was a member of the program in 1987 with a significant impact. There was also a young star, a young woman, who spent most of that

[22] *Selected Papers of Daniel Amit (1938–2007)*, Nicolas Brunel, Paolo del Giudice, Stefano Fusi, Giorgio Parisi and Misha Tsodyks Eds. (Singapore: World Scientific, 2013). https://doi.org/10.1142/8367
[23] Moshe Abeles : https://en.wikipedia.org/wiki/Moshe_Abeles

year with us. Her name was Elizabeth Gardner[24]. I don't know if you know the name.

FZ: Yes, we know her very well. We have been working these years on extending her ideas.

HG: [0:40:59] So let me tell you about Elizabeth Gardner. Elizabeth Gardner came to the Institut of Physique Théorique in Saclay, she worked with Bernard Derrida. She published two seminal papers[25]: one with him, and one alone. I should say something about her paper. Her paper was monumental in this field. It opened a whole new vista. Before I describe her work, and how it applied the replica method to this work, I must say a few things about her.

Elizabeth was a very modest and quiet young woman. We did not know that she had cancer before she joined us. She never complained. She had always a smile on her face. We were shocked to hear about her death a few months after she came home from Jerusalem. That was Elizabeth Gardner.

FZ: So you say that she spent several months in Jerusalem. We know that she already has cancer. But you didn't know if she got any treatment?

HG: [0:42:57] We did not know that she had cancer. I don't know if Bernard knew.

FZ: We already discussed with him and he said that he was not aware.

HG: No. None of us did. Nobody here was aware. By the way, Bernard also spent a few months with us here during that year.

Let me try to describe what was her main contribution. In a neural network there are two dynamics, two dynamical stages. One is the dynamics of learning. In that dynamics you change the J_{ij}, you adapt the J_{ij}s to converge to the desired attractors. Then, there is the other dynamics, the fast dynamics. This is the retrieval dynamics, when the network responds to an external stimulus. In the Hopfield model, we avoided the first step, and we guessed the J_{ij}. Instead of a long process of learning, we determined the synaptic efficacies by intelligent guesses.

---

[24] Elizabeth Gardner: https://en.wikipedia.org/wiki/Elizabeth_Gardner_(physicist)

[25] Elizabeth Gardner, "The space of interactions in neural network models," *J. Phys. A* **21**, 257 (1988). https://doi.org/10.1088/0305-4470/21/1/030; Elizabeth Gardner and Bernard Derrida, "Optimal storage properties of neural network models," *J. Phys. A* **21**, 271 (1988). https://doi.org/10.1088/0305-4470/21/1/031

Of course, in these cases, the J_{ij}s were all symmetrical because that is how we chose them. Our solution of the Hopfield model was based on a particular guess of these J_{ij}s.

Suppose you have a network with certain attractors and you want to learn, namely, to find the J_{ij}s that will ensure that the dynamics will lead you to one of those attractors. You can find them either by guessing, as we did in the Hopfield model, or by a certain learning procedure, which was known. Elizabeth showed that such a learning procedure, for this kind of a spin glass network, will converge to a solution, to a desired set of J_{ij}s, in a final number states, provided that such a solution exists. She could also estimate how long it will take to find such a solution. Then, she also showed that you can separate the question of finding those J_{ij}s from the questions of whether such J_{ij}s exist at all, and how many such sets there are. The J_{ij}s are the synaptic connections. They are continuous numbers, and there is a collection of N^2 such numbers connecting every two neurons in the network.

Supposed that you have certain prescribed attractors, and you want to know what J_{ij}s will give you a dynamics that will lead you to those attractors. If the number of such attractors is not too big there will be many such solutions. The assembly of such solutions will form a volume in a superspace of N^2 dimensions. She showed that this is a concave volume. As you increase the number of attractors that you want to embed in the system, that volume decreases until it decreases to zero and there is only one solution. This gives you the capacity limit. She found a way to calculate the volume of such solutions, under certain restrictions on the structure of the correlations between the embedded attractor configurations. Since she wanted a typical value, she had to calculate not the volume but the logarithm of the volume of this hyperspace of all the possible J_{ij} interactions. Again, in order to average over the disorder caused by the randomness of the attractor configurations, that required the replica symmetry method. Elizabeth demonstrated to us an elegant, beautiful derivation of her results.

PC: Can you tell us a bit more about your interactions with Elizabeth Gardner during that year, and how it eventually led to a publication with her?

HG: [0:49:24] During that year, we met daily, discussing over coffee and sandwich. She lived in Jerusalem. She gave seminars; we gave seminars. I had a student there, his name was Ido Yekutieli. He did his master's degree with me. I was looking for a problem, and I assigned him a problem which we started to discuss with her, applying her method to networks where we imposed *a priori* a certain level of asymmetry of the synaptic connections. Using her method one could impose certain restrictions on the J_{ij}s. For example, there were models, treated by her method, where the J_{ij}s were not continuous, but integers. In other models they were limited in size. After this long period of struggling with the question of symmetry and asymmetry, and what asymmetry does, we applied her method to a network in which we impose a certain amount of asymmetry. You can assign an asymmetry measure to the J_{ij}s . Essentially, it gives you a certain measure of the difference between the J_{ij} and J_{ji}. We could relate the storage capacity to that asymmetry measure. We discussed this idea at length with her. Actually, she liked that we kept her busy and involved. She died before this was completed. It was published in Journal of Physics A, a special issue dedicated to her memory[26]. By the way—I don't remember if I mentioned that—her paper was mentioned by the editors of the Journal of Physics A, where she used to publish—a British Journal—as one the 50 most influential papers in the history of that journal.

PC: Yes. Absolutely, without surprise. So this year was extremely successful, and from what I understand that led to the seeding of an Institute or Center at the Hebrew University. Can you tell us about that, or is there something else we should know about that year before moving on?

HG: [0:53:32] This was a very exciting year. You know [Pierre] Peretto?

PC: No, I don’t.

HG: He was a French scientist, working in Grenoble, who wrote a book about neural networks[27]. He also spent a year with us. Gérard Toulouse suggested that we invite him. Annette Zippelius[28], she also spent time with us. Of course, John Hopfield visited briefly. Miguel Virasoro was here at that time and a few others.

I know that if this interview would have taken place five years after that event, I would have much more to say. I remember this as a very exciting year. Now I am talking from memory and my memory fails to recall many more interesting details, except for the more basic elements, which I’ve already shared with you.

---

[26] E. Gardner, H. Gutfreund and I. Yekutieli, "The phase space of interactions in neural networks with definite symmetry," *J. Phys. A* **22**, 1995 (1989). https://doi.org/10.1088/0305-4470/22/12/005

[27] Pierre Peretto, *An Introduction to the Modeling of Neural Networks* (Cambridge: Cambridge Univesity Press, 1992). https://doi.org/10.1017/CBO9780511622793

[28] Annette Zippelius : https://en.wikipedia.org/wiki/Annette_Zippelius

FZ: Before we move on, I wanted to ask you something. You mentioned several times the interactions with École normale supérieure, but you didn't give us any details. How was it organized, the collaboration between your group and ENS?

HG: [0:55:15] All collaborations throughout my career were very personal. I spent several times mini-sabbaticals at École normale supérieure. When this activity developed, immediately after we did our work here, I spent a few months at Orsay and I gave a series of lectures—I even have notes of those lectures—on the recent developments of the statistical physics of neural networks. This was very well attended by students from the entire environment, from Saclay, from École normale supérieure, from many places. I spent some time at the École normale supérieure. I even have a joint paper from one of these visits with Marc Mézard, a Phys. Rev. Letters paper[29]. With Gérard Toulouse, we undertook a major initiative, that is, to edit a collection of seminal papers of those days from biology, computation, and physics, all of them around this field[30]. We wrote words of introduction to every one of these papers: biology and computation for physicists, physics for biologists, and so on. There were many, many things that happened in those days. They used to visit here, after that seminal year, several times. There was also one biologist in France that became interested in the development. That was Jean-Pierre Changeux[31]. He was also a kind of a remote partner, but we talked frequently. Many things happened in those years.

There were many international conferences. Every year, big conferences on statistical mechanics had a session on models of neural networks. I remember giving an invited talk in Sao Paulo on the topic "from statistical mechanics to neural networks and back"[32]. There was a lot of interest, there was a lot of activity.

And now to the final chapter. It's not the final chapter of the field, but the one that I was involved in, because my career ended—maybe too early—when I was drawn into university administration and science politics, when I was first elected rector of the Hebrew University and shortly after that

[29] H. Gutfreund and M. Mézard, "Processing of temporal sequences in neural networks," *Phys. Rev. Lett.* **61**, 235 (1988). https://doi.org/10.1103/PhysRevLett.61.235

[30] H. Gutfreund and G. Toulouse, *Biology and computation : a physicist's choice* (Singapore: World Scientific, 1994).

[31] Jean-Pierre Changeux : https://en.wikipedia.org/wiki/Jean-Pierre_Changeux

[32] STATPHYS 17 Workshop on Neural Networks and Spin Glasses, 8 – 11 August 1989 Porto Alegre, Brazil. See, e.g., *Neural Networks and Spin Glasses Proceedings of the STATPHYS 17 Workshop*, Walter K. Theumann Ed. (Singapore: World Scientific, 1990). https://doi.org/10.1142/0938

became its president. But in a way, that also played a role in the development of this activity at our university. I will tell you how.

We decided here that the activity in this field at our university would merit a special institute, a special program, a special PhD program that would bring together students from different disciplines, from physics, biology, psycholinguistics, computer science and so on. We submitted this proposal to the university. Initially, there was a strong opposition to that idea, mainly from the biologists. It was based on the argument that this was a waste of time. I remember that in those days, when we had one of those meetings, we invited Eric Kandel[33]. Eric Kandel visited us in Jerusalem, so we presented what we were doing and how excited we were. Some of our biological colleagues thought that we were arrogant. The truth is that we were also ignorant, and this proposal was a combination of both.

FZ: This is kind of typical of interdisciplinary interactions, right?

HG: [1:01:45] I remember Eric Kandel, I talked to him. I don't know if any of you knows Yiddish, but his reaction was, in Yiddish: "Thrown away money."[34] But he changed since then. I know that he changed his opinion. The presence in Columbia of Larry Abbott[35] and his group contributed to that.

I was the Rector of the university at that time, and I was convinced that this is what the university should do. As rector you cannot dictate, but you have ways to convince. Ultimately, we established this research institute, ICNC: Interdisciplinary Center for Neural Computation.[36] We witnessed the emergence of a new discipline and we were an important part of this process. I, particularly in my administrative position, my position at the top of the university, had on many occasions to describe and define what it means to be part of a new emerging field. One of the most successful projects of this Interdisciplinary Center for Neural Computation was its PhD programs. We had about 25-30 members from the four different disciplines that I mentioned. We had PhD students. Every year we had a retreat for a few days with the students. They did very different things but everybody knew what all the others were doing. The biologists knew what physicists were doing. To make them closely interacting we brought them to a retreat for a few days, a few nights. One rule of those retreats was that only the students lectured. After years of existence of ICNC, the Hebrew

[33] Eric Kandel : https://en.wikipedia.org/wiki/Eric_Kandel
[34] געלט אַרויסגעוואָרפֿענע, aroisgevorfene gelt : money thrown out.
[35] Larry Abbott: https://en.wikipedia.org/wiki/Larry_Abbott
[36] Interdisciplinary Center for Neural Computation: https://en.wikipedia.org/wiki/Interdisciplinary_Center_for_Neural_Computation

University brought in a peer-review committee headed by a renown medical scholar, Professor [Gerald] Fischbach[37]. Several Nobel laureates were [also] on that committee. The university received a very, very good report with specific recommendations on how to proceed. As a result of this report ICNC was transformed into a new center for brain research, ELSC, the Edmond and Lily Safra Center for Brain Sciences, with a new magnificent building designed by the illustrious architect, Norman Foster[38]. Come to Jerusalem to see it!

FZ: We have to visit it.

HG: [1:05:58] Long after I retired, I still continued to teach the basic course of the theory of neural networks, which now my younger colleagues have taken over. I look with great satisfaction and pride to where my younger colleagues are taking this initiative now.

PC: By curiosity, during those classes did you ever teach replica symmetry breaking ideas and concepts?

HG: This is a great question. I taught the results and I explained the trick. To go in detail through the calculation, every year when I got to that point I offered to the students who wanted to make an effort, a guided tour—guided by me—through a replica symmetry calculation with first-step replica symmetry breaking, line by line. I distributed my notes, and all those who volunteered met for a special meeting—not part of the course, but corollary to the course, and not part of the exam. All those who came to the course from physics and others from mathematics loved it. Actually, it was a great joy to take curious, intelligent students through certain *tour de force* step by step to surprising discoveries, to take them through such a route from the beginning to the final result.

PC: Wow. It's has been a really exciting tour. We only have one final question. Have you kept any papers, notes, correspondence from that epoch? If yes, do you have any plans to deposit them in an academic archive?

HG: [1:09:01] I drifted away from that field gradually, but then completely. I'm still officially a member of the ELSC[39], of this center. From time to time, I go to the seminars and lectures, trying to follow, not with great success. I'm now very deeply involved in something else. With another colleague,

[37] Gerald Fischbach : https://en.wikipedia.org/wiki/Gerald_Fischbach
[38] The Suzanne and Charles Goodman Brain Sciences Building by Norman Foster : https://en.wikipedia.org/wiki/Norman_Foster,_Baron_Foster_of_Thames_Bank
[39] Edmond & Lily Safra Center for Brain Sciences (ELSC)

[Jürgenn Renn,] I have already co-authored four books on different aspects of the genesis of general relativity and relativistic cosmology, how Einstein got to the equation of general relativity and so on[40]. I did not take the time even to organize my past. Your suggestion to engage in this interview, made me think. I will look at my notes and see what I can still salvage from there.

PC: That would be really wonderful. Very few people have kept those notes, so if you have it would be a great addition to the historical and intellectual record.

HG: [1:10:38] Can you tell me something about your initiative. How is it going?

PC: Maybe we'll stop the recording?

HG: Do you have any other questions?

PC: I think we've gone through everything. The only question I did not ask is if you had anything else to share. Otherwise, I think we've covered every point that we had identified ahead of time.

HG: That's it, so you can stop the recording.

PC: Thank you.

[40] Hanoch Gutfreund and Jürgen Renn, *The Formative Years of Relativity: The History and Meaning of Einstein's Princeton Lectures* (Princeton: Princeton University Press, 2017); *The Road to Relativity: The History and Meaning of Einstein's "The Foundation of General Relativity", Featuring the Original Manuscript of Einstein's Masterpiece* (Princeton: Princeton University Press, 2017); Albert Einstein, Hanoch Gutfreund and Jürgen Renn, *Relativity : the special & the general theory* (Princeton: Princeton University Press, 2019); Jürgen Renn and Hanoch Gutfreund, *Einstein on Einstein: Autobiographical and Scientific Reflections* (Princeton: Princeton University Press, 2020).