

The Glass Maze: Hiding Keys in Spin Glasses

Carlo A. Trugenberger

SwissScientific, chemin Diodati 10, CH-1223 Cologny, Switzerland
ca.trugenberger@bluewin.ch

Abstract: Key-binding mechanisms allow to use one's biometric features as a universal digital key without having them ever stored anywhere. Security, ease of use and privacy concerns are addressed in one stroke. The best known proposal for such a mechanism, the Fingerprint Vault, treats biometric data as projections of a polynomial encoding the key. Its security is based on the difficulty of polynomial reconstruction. Here I propose a new key-binding mechanism based on associative pattern recall and making use of a totally different security principle, that of the difficulty of energy optimization of spin glasses. The idea is to exploit the mixed ferromagnetic and spin glass phase of the Hopfield neural network to encode the key as a local minimum configuration of the energy functional, "lost" amidst the exponentially growing number of valleys and minima representing the spin glass. The correct fingerprint will be able to retrieve the key by dynamical evolution to the nearest attractor. Other fingerprints will be driven far away from the key. Known vulnerabilities of the Fingerprint Vault are eliminated by this new security principle.

1 Introduction: Biometric Cryptosystems

Traditional cryptography uses keys to encipher messages. Achieving a high degree of security with this approach requires that keys are long and random, which makes them also difficult to memorize, maintain and share. In addition, every service typically wants to issue its own key, thereby compounding the problem for end customers. Finally, the authentication process is based on keys and not on users and cannot thus provide a basis for non-repudiation.

All these problems can be solved by biometric cryptosystems [UI04], in which cryptography and biometrics are combined. Cryptography provides the necessary security, while biometrics ties this security to a specific user: anatomical traits cannot be lost or forgotten, are difficult to copy and share, can be used as a unique hub for multiple keys related to different services and provide non-repudiation. The resulting systems benefit thus from the strengths of both fields.

In biometric cryptosystems, a cryptographic key is combined with the biometric template of a user in such a way that the key cannot be revealed without a concomitant successful biometric authentication. Biometric cryptosystems can operate in one of the following three modes: *key release*, *key binding* or *key generation*.

In the *key release* mode, biometric authentication is completely decoupled from the key

release mechanism. The biometric template and the key are both stored in a central repository, but they constitute completely separate entities: the key is released only if and after the biometric matching is successful. Because of its simplicity this is the most common approach. Simplicity, however comes at the price of two weaknesses. First, the biometric template is not secure: template security is a critical issue in biometric systems because stolen templates cannot be revoked. Secondly, since authentication and key release are decoupled, it is possible to override the biometric matcher using a Trojan horse program.

In the *key binding* mode, the key and the template are monolithically bound within a cryptographic framework. It is impossible to decode the key without any knowledge of the user's biometric data. An appropriately designed matching algorithm is used to perform authentication and key release in a single step, thus removing any possibility of Trojan horse attacks. Since no "naked" template is ever stored, the privacy issue is also solved.

In the *key generation* mode, neither the key nor the template are stored in any form. The key is derived directly from the live biometric data provided by a user.

Though it is easy to implement a biometric cryptosystem in the *key release* mode, such a system is not appropriate for high security applications because of its major vulnerabilities. Biometric cryptosystems that work in the *key binding/generation* modes are much more secure but difficult to implement due to large intra-class variations in biometric data.

One of the best studied and well accepted key binding approaches is the Fingerprint Vault [CKL03], which specializes to fingerprint minutiae the generic Fuzzy Vault scheme [JS02]. The idea is to encode a key in the coefficients of a given polynomial P , use the abscissa x_i of fingerprint minutiae to generate the set T of points $(x_i, P(x_i))$ and add random chaff points forming the set F . The union V of T and F constitutes the fingerprint vault. The principle of the vault is that only the genuine fingerprint will be able to separate the true set T from the chaff points F and thereby reconstruct the polynomial P , and thus the key, by suitable error correcting codes.

The Fingerprint Vault makes use only of the position of fingerprint minutiae for the vault locking set. Intra-class variance is dealt with by computing average minutiae locations over multiple scans at enrollment and quantizing positions by using pixel coordinates. Chaff points have to be placed outside regions determined by the statistical variance of minutiae locations.

The security of the Fuzzy Vault idea is based on the difficulty of polynomial reconstruction. While the idea is brilliant, it has been claimed [Mi07] that it is vulnerable to brute force attacks (at least in its original formulation based on the data of one single fingerprint), and to cross-matching, surreptitious key inversion and blended substitution attacks [SB07]. It is therefore crucial to find also alternative key binding schemes which lack these vulnerabilities. In this paper I propose such a completely new scheme, which is based on a totally different security principle, that of the difficulty of optimizing the energy functional of a spin glass [MPV87].

2 Spin Glasses

Spin glasses [MPV87] are networks of spins, i.e. binary variables, with symmetric random (positive and negative) interactions. They are characterized by an energy function

$$E = -\frac{1}{2} \sum_{i \neq j} w_{ij} s_i s_j, \quad s_i = \pm 1, \quad i, j = 1 \dots N. \quad (1)$$

The presence of both positive and negative random interactions w_{ij} leads to strong frustration. This means that, in any configuration, there are many triplets of spins for which it is not possible to minimize the energy of all pairwise bonds at the same time. As a consequence, the energy landscape of spin glasses becomes extremely complex, with a number of local minima growing exponentially with the number N of spins [Ta80, TE80]. Finding energy minima of a spin glass is a "hard" computational problem: from the point of view of complexity theory this problem is NP complete.

Given the exponential number of local minima, the new idea is to hide the key as the spin configuration corresponding to one specific minimum of the energy landscape of a spin glass. Finding the exact position of this minimum with no a priori information on its approximate location is a NP complete problem. All known algorithms to solve this problem are highly inefficient and require at least super-polynomial, if not exponential, running time in system size. Typical combinatorial optimization heuristics, such as simulated annealing, branch and bound or evolutionary techniques cannot help, they require only polynomial computing resources but they provide only approximate solutions for the ground state [MPV87, PIM06, Ba91]. While approximate solutions may be sufficient for engineering problems, they are of no use in the present cryptographic framework. Moreover, the key will typically not correspond to the global minimum of the energy landscape but only to a local minimum: this excludes also the recently developed quantum adiabatic optimization techniques [Ka11]. Exhaustive search is essentially the only attack option and choosing a sufficiently large N should guarantee the cryptographic security of the model.

Not any spin glass will do, however. If the generic complexity of such systems guarantees the cryptographic security of hidden keys, it will also prevent the legitimate retrieval of the key by its owner. The natural retrieval mechanism is a hill-descent dynamics starting from an initial configuration close enough to the key, a configuration provided by the key owner, e.g. the position of his fingerprint minutiae. The system should allow the efficient associative retrieval of the key from a similar (albeit different) configuration while still be complex enough to render any guessing attempt hopeless. In statistical mechanics terms, an efficient associative retrieval of information corresponds to a ferromagnetic phase of the system. Therefore, what is needed is a system with mixed ferromagnetic and spin glass phases. Fortunately, exactly such a system exists: it is the Hopfield neural network model of associative memory [MR90]. The new proposal for a key-binding scheme presented here is based on this model or variants thereof.

3 Neural Networks for Associative Memory: the Hopfield Model

Historically, the interest in neural networks [MR90] has been driven by the desire to build machines capable of performing tasks for which the sequential circuit model of Babbage and von Neumann are not well suited, like pattern recognition, categorization and generalization. The Hopfield model is one of the best studied and most successful neural networks. It was designed to model one particular higher cognitive function of the human brain, that of associative pattern retrieval or associative memory.

Contrary to traditional computers, which require a lookup table (RAM), biological information retrieval is possible on the basis of a partial knowledge of the memory content, without knowing an exact storage location: noisy or incomplete inputs are permitted.

The Hopfield model consists in an assembly of N binary neurons s_i , $i = 1 \dots N$, which can take the values ± 1 representing their firing (+1) and resting (-1) states. The neurons are fully connected by symmetric synapses with coupling strengths $w_{ij} = w_{ji}$ ($w_{ii} = 0$). Depending on the signs of these synaptic strengths, the couplings will be excitatory (> 0) or inhibitory (< 0). The dynamical evolution of the network state is defined by the random sequential updating (in time t) of the neurons according to the rule

$$s_i(t+1) = \text{sign}(h_i(t)) , \quad (2)$$

$$h_i(t) = \sum_{j \neq i} w_{ij} s_j(t) , \quad (3)$$

where h_i is called the local magnetization.

The synaptic coupling strengths are chosen according to the Hebb rule

$$w_{ij} = \frac{1}{N} \sum_{\mu=1 \dots p} \sigma_i^\mu \sigma_j^\mu , \quad (4)$$

where σ_i^μ , $\mu = 1 \dots p$ are p binary patterns to be memorized. An associative memory is now defined as a dynamical mechanism that, upon preparing the network in an initial state s_i^0 retrieves the stored pattern σ_i^λ that most closely resembles the presented pattern s_i^0 , where resemblance is determined by minimizing the Hamming distance, i.e. the total number of different bits in the two patterns. As emerges clearly from this definition, all the memory information in a Hopfield neural network is encoded in the synaptic strengths, which correspond to the spin interactions of the spin glass model above.

It can be easily shown that the dynamical evolution (2) of the Hopfield model satisfies exactly this requirement for an associative memory. This is because:

- The dynamical evolution (2) minimizes the energy functional (1), i.e. this energy functional never increases when the network state is updated according to the evolution rule (2). Since the energy functional is bounded by below, this implies that the network dynamics must eventually reach a stationary point corresponding to a, possibly local, minimum of the energy functional.

- The stored patterns σ_i^μ correspond to, possibly local, minima of the energy functional. This implies that the stored patterns are attractors for the network dynamics (2). An initial pattern will evolve till it overlaps with the closest (in Hamming distance) stored pattern, after which it will not change anymore.

Actually the second of these statements must be qualified. Indeed, the detailed behavior of the Hopfield model depends crucially upon the loading factor $\alpha = p/N$, the ratio between the number of stored memories and the number of available bits. This is best analyzed in the thermodynamic limit $p \rightarrow \infty$, $N \rightarrow \infty$, in which the different regimes can be studied by statistical mechanics techniques and characterized formally by the values of critical parameters.

For $\alpha < \alpha_{1\text{crit}} \simeq 0.051$, the system is in a ferromagnetic (F) phase in which there are global energy minima corresponding to all stored memories. The former differ from the original input memories only in a few percent of the total number of bits. Mixing between patterns leads to spurious local energy minima. These, however are destabilized at sufficiently high temperatures (see below) and thus an exhaustive search for all stored memories can be efficiently organized by optimization heuristics such as simulated annealing.

For $\alpha > \alpha_{2\text{crit}} \simeq 0.138$, the system is in a spin glass (SG) phase in which all retrieval capabilities are lost due to an uncontrolled proliferation of spurious memories.

For $\alpha_{1\text{crit}} \simeq 0.051 < \alpha < \alpha_{2\text{crit}} \simeq 0.138$ the system is in a mixed spin glass and ferromagnetic phase. There are still minima of sizable overlap with the original memories but they are now only metastable states. The true ground state is the spin glass, with an exponentially increasing number of minima due to the mixing of original memories. The spin glass phase is orthogonal to all stored memories. If an input pattern is sufficiently near (in Hamming distance) to one of the original memories it will be trapped by the corresponding metastable state and the retrieval procedure is successful. On the other hand, if the input pattern is not sufficiently close to one of the stored memories, the network is confused and it will end up in a state very far from all original memories. In this phase, pattern retrieval is still efficient while it is computationally hard to perform exhaustive searches for all energy minima. This is exactly the property needed to construct a cryptographically secure fuzzy key retrieval mechanism.

4 The Glass Maze

The spin glass phase limiting the loading factor of the Hopfield model is its main drawback as an associative memory model of the human brain, its design application. It is, however an ideal feature for hiding information : exploiting the disorder limiting the intended scope of the model to one's advantage one can turn the Hopfield model into a powerful cryptographic tool. Here is how I propose to do so.

I will henceforth specialize to fingerprints, although the model is applicable, mutatis mutandis also to other biometric traits. Fingerprints [Bo04] can be uniquely characterized

by features of their print pattern. One of the most popular such set of features are fingerprint minutiae, of which there are three types : ridge endings, bifurcations and short ridges. Minutiae are typically characterized by their pixel coordinates (x_i, y_i) on the two-dimensional plane and by a direction unit vector \mathbf{r}_i (or equivalently by an angle ϕ_i with respect to a fixed axe). As in the original Fingerprint Vault [CKL03], I will consider only minutiae positions as the relevant variables, although the whole model can be easily extended to include minutiae directions.

Each fingerprint is thus characterized by a variable set of M minutiae points with pixel coordinates (x_i, y_i) , $i = 1 \dots M$, on the fingerprint image. The first step in setting up the model consists in choosing an appropriate quantization of the fingerprint image in N squares. Each such square will be considered as a neuron in a Hopfield model, and this neuron takes the value $+1$ if a minutiae point is contained within the square and -1 otherwise. This way, a fingerprint becomes one particular state configuration σ_i^{fp} in a Hopfield model.

Intra-class variance, however, invariably adds noise to minutiae coordinates. When comparing multiple scans of the same finger, two types of variability of feature coordinates can be established : global, rigid transformations due to translations and/or rotations of the finger (I neglect scale transformations assuming that the same scanner is used for the enrolling and key release processes) and additional local deformations due to elastic non-linear deformations. The magnitude of both has been analyzed by [UJ04, UJ06, UPJ05] in data from the IBM GTDB database, images of 300 X 400 pixels with 500 DPI resolution. The average coordinate difference of the same features in mated fingerprint pairs due to global transformations is of about 20 pixels, with values as large as 45 pixels occurring with non negligible probability 0.09. After fingerprints have been aligned by eliminating translational and rotational transformations, a residual average coordinate difference of 4 pixels has been detected : this corresponds to the local noise component.

Neither global nor local transformations should pose a problem for the new model if the quantization is chosen appropriately. Indeed, the key-binding scheme is, by construction, an associative memory, capable to reconstruct and retrieve also noisy and corrupted inputs. A certain degree of intra-class variability is automatically tolerated by the associative nature of the Hopfield model. For example, choosing $N=256$, the horizontal and vertical pixels are aggregated in 16 groups that, for the IBM GTDB database would contain 19 or 25 pixels each. This is exactly the average variability due to global transformations [UJ04, UJ06, UPJ05]. The residual average local deformations are well below the dimension of the quantization window. With such a choice of parameters, it is to be expected that local deformations are eliminated by the quantization, whereas global deformations are corrected by the associative pattern recognition. Only experiments, however, can really decide if an alignment (elimination of global transformations) is needed before the Hopfield dynamical retrieval for a given quantization window. This can be a standard alignment procedure or a dynamical alignment. Indeed, contrary to the fingerprint vault, alignment can be obtained as a self-organizing, dynamical process within the model (see below).

At enrollment, a fingerprint is assigned to a particular neuron configuration as described above. The key (for example a 256-bit key) to be bound to this fingerprint is now chosen

by modifying randomly a certain number k of components of this fingerprint neuron configuration. Let us call the resulting neuron configuration σ_i^{key} : this is the reference to be stored for comparison. The Hopfield model realizing the desired key-binding mechanism is then defined by the interaction parameters:

$$w_{ij} = \frac{1}{N} \sum_{\mu=1 \dots p} \sigma_i^\mu \sigma_j^\mu, \quad (5)$$

$$\sigma_i^1 = \sigma_i^{\text{key}}, \quad (6)$$

$$\sigma_i^\mu = \text{random patterns, } \mu = 2 \dots p. \quad (7)$$

The total number p of patterns must be chosen in such a way that the loading factor $\alpha = p/N$ satisfies the relation $\alpha 1_{\text{crit}} \simeq 0.051 < \alpha < \alpha 2_{\text{crit}} \simeq 0.138$, so that the addition of random noise to the key pattern creates a spin glass phase mixed in with the ferromagnetic recall phase. A loading factor near 0.1 is e.g. a good choice. Notice that the addition of random noise at this loading factor typically modifies the desired key pattern in a few percent of its bits [MR90]. The actual key corresponds to this modified pattern and it must be checked that the original fingerprint pattern still lies in the basin of attraction of the true key. Otherwise the number of modified bits k has to be lowered. The synaptic coupling strengths w_{ij} are the data needed by the retrieval algorithm.

When a live fingerprint is presented to the system, its minutiae coordinates are extracted and matched to a corresponding initial neuron configuration s_i^0 . Starting from this initial configuration, the network is evolved according to the Hopfield dynamics (2). Under this evolution, the neuron configuration corresponding to the fingerprint evolves towards one of the many minima of the Hopfield energy landscape. If the live fingerprint is the correct one, the fixed point will correspond to the key. If another fingerprint is presented, which differs enough from the correct one, the evolution will drive the network towards a completely different energy minimum and the key will not be retrieved.

5 Performance and Security

The False Accept Rate (FAR) and False Reject Rate (FRR) of the proposed system depend on the two model parameters N and α . Choosing a coarse quantization grid (low N) will lower the FRR, since the quantization window will tend to encompass all coordinate variations due to local noise. Lowering N will also increase the FAR, since the probability that two different fingerprints are matched to very near neuron configurations is enhanced. Increasing N , on the contrary, will increase the FRR and lower the FAR. Detailed tests are necessary to establish a Receiver Operating Characteristic (ROC) curve that plots the genuine accept rate (GAR) against the FAR. This curve can be established by varying N at fixed loading α .

FAR and FRR are also influenced by the size of the basin of attraction of the key pattern and this depends on the loading factor α . The larger this size, the more local variations in the presented fingerprint minutiae features are tolerated, which decreases the FRR. On the

other side a larger basin of attraction increases the probability that a genuinely different fingerprint configuration is evolved to the same key, i.e. it increases the FAR. The average radius of the basins of attraction can be estimated by statistical mechanics methods [ZC89]. In the thermodynamic limit it depends only on the loading factor $\alpha = p/N$ and permits approximately 10% of wrong bits in the limit $\alpha \rightarrow 0$, while approaching zero when $\alpha \rightarrow \alpha_{2\text{crit}} \simeq 0.138$. Numerical studies in finite size samples of some thousand bits [ZC89, YD09], however, have revealed considerably larger basins of attraction. The best value of α for the model has to be established by testing on realistic finite-size samples.

The number N of neurons and the loading factor α determine also the model security complexity. Two different measures for the model security can be established. The first is the difficulty of FAR attacks, in which random fingerprints are sequentially submitted for key release. These are brute force attacks, in which the whole configuration space is sequentially explored in search of an input that evolves to the key. Relevant for this security measure is the size of the basins of attraction. A realistic basin of attraction permitting to retrieve the key with less than 5% different bits in the input, would contain, in our example $N=256$, about 2.5×10^{21} configurations. The total number of possible configurations is $2^{256} \simeq 1.16 \times 10^{77}$. Of these not all are realistic fingerprints, though. Taking into account that typical fingerprints have 20-40 minutiae, one can estimate that the quantization implies that realistic configurations have a maximum of, say, 35 positive neurons. Of these there are about 2×10^{43} . This gives a probability of about $1.25 \times 10^{-22} \simeq 2^{-73}$ of retrieving the key from a random input.

Instead of a blind search through all configuration space, one can try to crack the model by optimization heuristics that are able to find minima in polynomial time in system size N . To this end one must explore, on average, all the minima, local and global in the energy landscape. A second measure of the model security complexity is thus given by the average number of such minima at a given N . This can be estimated as follows. The spin glass phase of the Hopfield model is in the same universality class as the Sherrington-Kirkpatrick model [KS78]. The average number of minima for this spin glass model has been computed [Ta80, TE80] as $\exp(0.2 \times N)$. While this result is, strictly speaking, valid only in the thermodynamic limit $N \rightarrow \infty$, it can be used as an analytical estimate of the ground state complexity. In our example $N = 256$ it would imply $1.7 \times 10^{22} \simeq 2^{74}$ minima in the spin glass phase, a figure remarkably close to the previously computed security against brute force FAR attacks. Note also that, since the key memory is metastable, all attack algorithms based on thermal noise would never find the key : after a certain time spent around the key configuration they would inevitably get confused and driven far away from it. The security complexities 2^{73} and 2^{74} for a 256-bit key have to be compared with the security complexity 2^{69} for a 128-bit key of the Fingerprint Vault [CKL03].

A definite advantage of the present model over the Fingerprint Vault is its security against cross-matching attacks. This arises because, contrary to the Fingerprint Vault [SB07], different implementations of the glass maze for the same fingerprint do not share any model data. Disorder is not simply added to fixed data but templates are actually embedded non-linearly in an energy landscape which is both inaccessible to exhaustive analysis and chaotic [Par02].

Furthermore, the Glass Maze, contrary to the Fingerprint Vault [SB07], is resistant against

surreptitious key inversion attacks. Even if the key is stolen, there are still about 2.5×10^{21} (in the example above) configurations in the key basin of attraction that could all correspond to the original fingerprint template. The probability of guessing the biometric data from a stolen key is thus negligibly small.

Finally, the Glass Maze is also secure against blended substitution attacks, another weakness of the Fingerprint Vault [SB07]. There is simply no systematic way to insert an attacker's fingerprint template into the same basin of attraction of a legitimate user without altering severely the chaotic energy landscape and thus also the key.

I would like to conclude this section by pointing out that the model storage requirement is also determined by the neuron number N and the loading factor α . Let me compute it explicitly for the example of a 256-bit key and a loading factor $\alpha = 0.1$. Each interaction term w_{ij} is a number in the interval $[-p, +p]$ (the factor $1/N$ is not important here). For $N=256$ and $\alpha = 0.1$ one needs 6 bits to represent such a number. Since there are a total of $256^2 = 65'536$ of them in the definition of the corresponding model one obtains a total storage requirement of around 50 kbyte, without taking into account possible space saving by compression. This is about 100 times more than the storage requirement for a standard fingerprint template. Note however, that multiple fingerprints can easily be encoded in the same network by using their neuron configurations as memories : each fingerprint configuration will act as noise for all others. In our example $N=256$, up to about 25 fingerprints can be encoded, which would give a storage requirement of about 2 kbyte per fingerprint without sensibly altering the security complexities.

6 Dynamical Alignment

Alignment is the process of elimination of rigid rotations and translations between enrolled and live fingerprint templates prior to comparison (again I neglect scale transformations assuming that the same device is used for enrollment and key release). As explained above, the choice of an appropriate quantization is probably sufficient to eliminate the need for pre-alignment in the Glass Maze model, since local deformations fall well within the quantization window and typical rigid, global transformations are compatible with the natural error-correcting nature of the associative retrieval process.

For small quantization grids, instead, a pre-alignment step is needed and two possibilities are available. Either one resorts to standard alignment techniques or one exploits the nature of the model itself. Indeed, contrary to other biometric authentication systems, in which special alignment techniques are unavoidable, the present model contains in itself a natural mechanism for dynamically aligning minutiae prior to key search.

In order to explain this mechanism I have to introduce first two generalizations of the Hopfield model. The first concerns stochastic neurons and is realized by turning the deterministic evolution law (2) into a stochastic law by adding thermal noise :

$$\text{Prob}[s_i(t+1) = +1] = f[h_i(t)] , \quad (8)$$

where the activation function f is the Fermi function

$$f(h) = \frac{1}{1 + \exp(-2\beta h)} , \quad (9)$$

and the parameter β plays thus the role of an inverse temperature, $\beta = 1/T$. This has to be understood as a fictitious temperature for the evolution law and not as the physical temperature at which the network operates. In the limit of zero temperature, $\beta \rightarrow \infty$, the Fermi function approaches the unit step function and the stochastic neural network goes over into the original deterministic one.

In a deterministic network, neurons are either permanently active or permanently dormant, depending on the sign of the local magnetization field h . In a stochastic network, the neuron activities fluctuate due to thermal noise. Even for positive local magnetization h , there is a non-vanishing probability that the neuron will flip in the next evolution step :

$$\text{Prob}[s_i(t+1) = s_i(t)] = \frac{\exp[\beta h_i(t)s_i(t)]}{2 \cosh[\beta h_i(t)s_i(t)]} , \quad (10)$$

$$\text{Prob}[s_i(t+1) = -s_i(t)] = \frac{\exp[-\beta h_i(t)s_i(t)]}{2 \cosh[\beta h_i(t)s_i(t)]} . \quad (11)$$

As a consequence, the network acquires a non-vanishing probability of jumping out of a local energy minimum.

The second generalization we need consists of non-vanishing thresholds θ_i in the local magnetization,

$$h_i(t) = \sum_{j \neq i} w_{ij} s_j(t) + \theta_i , \quad (12)$$

These play the role of an external magnetic field acting on the spins.

To dynamically align a live fingerprint one must construct an algorithm that matches the neuron configuration s_i^0 of its minutiae to a configuration as close as possible to the original minutiae configuration σ_i^{fp} of the enrolled fingerprint. I shall assume here that the same device is used for enrollment and key release, so that scale transformations can be neglected and the two neuron configurations s_i^0 and σ_i^{fp} differ only by global rotations and translations.

Following an idea of Dotsenko [Do88] I shall consider a stochastic Hopfield model with thresholds determined by the neuron configuration s_i^0 corresponding to the live fingerprint. For ease of presentation I shall henceforth change notation and label neurons not by a sequential integer index i , but rather by the discrete coordinate vectors \mathbf{r} of the square lattice defined by the quantization of the two-dimensional plane of the print, $s_i(t) \rightarrow s(\mathbf{r}, t)$. The thresholds are thus given by

$$\theta(\mathbf{r}, t) = h_0 s^0(\mathbf{r}, t) , \quad (13)$$

and describe an external magnetic field proportional to the live fingerprint configuration : the parameter h_0 describes the strength of this magnetic field.

The idea is to introduce rotations $\phi(t)$ and discrete translations $\mathbf{a}(t)$ of the square lattice as dynamical, slow variables and to minimize the energy

$$E = -\frac{1}{2} \sum_{\mathbf{r} \neq \mathbf{r}'} w_{\mathbf{r}\mathbf{r}'} s(\mathbf{r})s(\mathbf{r}') - \sum_{\mathbf{r}} h_0 s^0([\phi(t)\mathbf{r}] + \mathbf{a}(t)) s(\mathbf{r}), \quad (14)$$

of a generalized Hopfield model with external magnetic field with respect to these slow variables. Here, $[\dots]$ denotes the whole part, so that the arguments of s^0 are always discrete lattice vectors.

To this end one imposes, e.g. periodic boundary conditions on the lattice and one chooses initial values $\phi(t=0)$ and $\mathbf{a}(t=0)$ for the rigid transformations. Then the stochastic neurons are thermalized according to the fast dynamics (8). Finally, the slow variables are evolved according to the standard steepest descent relaxation equations

$$\phi(t+1) - \phi(t) = -\lambda_\phi \partial E / \partial \phi + \zeta(t), \quad (15)$$

$$\mathbf{a}(t+1) - \mathbf{a}(t) = -\lambda_{\mathbf{a}} \delta E / \delta \mathbf{a} + \eta(t), \quad (16)$$

where λ_ϕ , $\lambda_{\mathbf{a}}$ denote the descent steps, ζ and η are possible ordinary temperature noise terms and δ denotes finite differences. Then, the whole process is repeated : neurons are thermalized and the global transformation parameters are changed so that the total energy is decreased again.

The important point is that, as long as all overlaps

$$M^\mu = \frac{1}{N} \sum_{\mathbf{r}} \sigma^\mu(\mathbf{r}) s^0([\phi\mathbf{r}] + \mathbf{a}), \quad (17)$$

with the stored memories are small, $M^\mu \ll h_0$, the external magnetic field in (14) dominates the thermalization dynamics and, as a consequence, the neurons follow rigidly the external field

$$\langle s(\mathbf{r}, t) \rangle = \tanh(\beta h_0) s^0([\phi(t)\mathbf{r}] + \mathbf{a}(t)), \quad (18)$$

Therefore, at this stage, the search for the correct pattern is a wandering over neuron configurations corresponding to rigid transformations of the initial live fingerprint configuration: this is what I call dynamical alignment.

Dotsenko [Do88] has shown that, if the following conditions are satisfied:

- For some ϕ^* and \mathbf{a}^* the pattern $s^0([\phi^*\mathbf{r}] + \mathbf{a}^*)$ has a finite overlap $M^{\text{key}} = O(1)$ with one of the stored memories, in our case the key configuration, and has no finite overlaps with any other memory, $M^\mu = O(1/N^2)$ for $\mu \neq \text{key}$;
- All stored memories (and the initial fingerprint configuration) have a finite spatial correlation length R_c , $(1/N) \sum_{\mathbf{r}} \sigma^\mu(\mathbf{r}) \sigma^\mu(\mathbf{r} + \mathbf{R}) \propto \exp(-|\mathbf{R}|/R_c)$,
- $1/N < h_0 < 1$

- The temperature T lies in the range $1 > T > T_c$ with T_c given by the solution of the transcendental equation $T_c/\tanh^2(h_0/T_c) = 1$,

then the network will localize efficiently around the configuration σ^{key} , i.e. the steepest descent will drive the network toward σ^{key} until the overlap $M^{\text{key}} \simeq h_0$.

Operating the network at a finite temperature T in the range above guarantees that the evolution does not get trapped around an undesired minimum while still guaranteeing the convergence toward the key pattern. Since, in this stage, the search is limited to rotations and translations of the live fingerprint configuration, the result must be very close to the enrolled fingerprint configuration σ^{fp} , the only remaining difference is due to possible local variability. In any case, this process should guarantee that an initial configuration corresponding to a correct fingerprint is brought into the basin of attraction of the key. After this dynamical alignment, the key can be retrieved by the standard Hopfield dynamics described previously. In this final step the search takes place over independent variations of all neurons till $M^{\text{key}} \simeq h_0 \rightarrow M^{\text{key}} = 1$.

7 Conclusion

In conclusion, I have shown how the Hopfield model of associative memory can be used as an efficient key-binding biometric cryptosystem. Its pattern-retrieving capabilities are exploited to bind a fingerprint minutiae configuration to a corresponding key configuration which represents a specific minimum in an energy landscape. The spin glass phase limiting the original application as a biological information storage model, instead, provides the disorder and complexity necessary to hide the key configuration in a "maze" of valleys and minima so that the resulting cryptosystem is robust against attacks.

Detailed tests are necessary to evaluate the performance and security of this new model and, in particular, to find optimal parameters for concrete applications.

References

- [Ba91] Banzhaf, W.: Finding the Global Minimum of a Low-Dimensional Spin-Glass Model. In (Becker, J. D.; Eisele, I.; Mündemann, F. D. eds.): *Parallelism, Learning, Evolution, Lectures Notes in Artificial Intelligence*, Springer Verlag, Berlin, 1991.
- [Bo04] For a review see: Bolle, R. M.; Connell, J. H.; Pankanti, S.; Ratha N. K.; Senior A. W. : *Guide to Biometrics*, Springer Verlag, Berlin, 2004.
- [CKL03] Clancy, T. C.; Kiyavash, N.; Lin, D. J. : Secure Smartcard-Based Fingerprint Authentication. In: *Proceedings of the ACM SIGMM Workshop on Biometric Methods and Applications*, 2003, pp. 45-52.
- [Do88] Dotsenko, V. S.: Neural Networks: Translation-, Rotation- and Scale-Invariant Pattern Recognition, *J. Phys. A: Math. Gen.* 21, 1988, p. L783.

- [JS02] Juels A.; Sudan, M. : A Fuzzy Vault Scheme. In (Lapidith, A.; Teletar, E. eds.): Proceedings of the IEEE International Symposium on Information Theory, 2002, p. 408.
- [Ka11] Karimi K.; Dickson N. G.; Hamze F.; Amin M. H. S.; Drew-Brook M.; Chudak F. A.; Bunyk P. I.; Macready W. G.; Rose G. : Investigating the Performance of an Adiabatic Quantum Optimization Processor, arXiv:1006.4147v4, 2011.
- [KS78] Kirkpatrick, S.; Sherrington, D. : Infinite-Ranged Models of Spin Glasses, Phys. Rev. B17, 1978, pp. 4384-4403.
- [MPV87] For a review see: Mezard, M.; Parisi, G.; Virasoro, M. A. : Spin Glass Theory and Beyond, World Scientific, Singapore, 1987.
- [Mi07] Mihailescu, P. : The Fuzzy Vault for Fingerprints is Vulnerable to Brute Force Attacks, eprint arXiv:0708.2974.
- [MR90] For a review see: Müller, M.; Reinhard, J. : Neural Networks, Springer Verlag, Berlin, 1990.
- [Par02] For a review see Pardo M. S. : Large Scale Excitations in Disordered Systems, Doctoral Thesis, University of Barcelona, 2002.
- [PIM06] Percus, A.; Istrate, G.; Moore, C. : Computational Complexity and Statistical Physics, Oxford University Press, Oxford, 2006.
- [SB07] Scheirer, W. J.; Boulton, T. E. : Cracking Fuzzy Vaults and Biometric Encryption. In: Proceedings of the Biometric Symposium, 2007, pp. 1-6.
- [Ta80] Tanaka, F. : Ground-State Properties of the Infinite-Range Model of a Spin Glass, J. Phys. C13, 1980, pp. 1951-1955.
- [TE80] Tanaka F.; Edwards, S. F. : Analytic Theory of Ground State Properties of a Spin Glass: I. Ising Spin Glass, J. Phys. F10, 1980, pp. 2769-2778.
- [UJ04] Uludag U.; Jain, A. : Fuzzy Vault for Fingerprints. In: Proceedings of the Workshop on Biometrics: Challenges Arising from Theory and Practice, 2004, pp. 13-16.
- [UJ06] Uludag U.; Jain, A. : Securing Fingerprint Templates: Fuzzy Vault with Helper Data. In: Proceedings of the IEEE Workshop on Privacy Research in Vision, New York, 2006.
- [UPJ05] Uludag, U.; Pankanti S.; Jain, A. : Fuzzy Vault for Fingerprints. In: Proceedings of AVBPA 2005 Workshop, Lecture Notes in Computer Science, Springer Verlag, Berlin, 2005, pp. 310-319.
- [UI04] Uludag, U.; Pankanti S.; Prabhakar S.; Jain, A. : Biometric Cryptosystems: Issues and Challenges. In: Proceedings of the IEEE International Workshop, 2004, pp. 948-960.
- [YV05] Yang S.; Verbauwhede I. : Automatic Secure Fingerprint Verification System Based on Fuzzy Vault Scheme. In: Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, 2005, pp. 609-612.
- [YD09] Yasar F.; Dilaver M. : The Systematic Simulation of the Hopfield Model by Multicanonical Algorithm, Chinese Journal of Physics 47, 2009, pp. 226-237.
- [ZC89] Zagrebnev V. A.; Chyrkov, A. S. : The Little-Hopfield Model: Recurrence Relations for Retrieval Pattern Errors, Sov. Phys. JETP 68, 1989, p. 153.

