Teaching the Mathematical Foundations of Lattice-Based Cryptography

Wei Gao, Tian Li

Department of Computer Science, North China Electric Power University (Baoding), Baoding, China

Abstract

With the development of post-quantum cryptography, information security students are required to master the mathematical foundations of lattice-based cryptography. However, the current course Foundations of Information Security Mathematics still follows the traditional sequence of number theory and abstract algebra and lacks a systematic, vector-space-oriented presentation of lattice structures. To address this gap, this paper constructs a knowledge framework covering linear spaces, Euclidean metric structures, finite modular algebraic structures, and lattice-based hard problems, and analyzes the key cognitive transitions involved in learning. A teaching path centered on cognitive transition is proposed, in which geometric visualization, representation comparison, numerical experiments, and structural correspondence are used to support students' structural understanding. In this way, SVP and CVP are unified as metric optimization problems, and LWE is interpreted as a noisy linear structure in a finite modular discrete linear space.

Keywords

Foundations of Information Security Mathematics; post-quantum cryptography; lattice-based cryptography; linear space.

1. The Necessity of Introducing the Mathematical Foundations of Lattice-Based Cryptography

With the rapid progress of quantum computing, traditional public-key cryptosystems based on the hardness of integer factorization and the discrete logarithm problem face potential threats. Constructing post-quantum cryptographic systems capable of resisting quantum attacks has therefore become an important research direction in information security. The standardization results of post-quantum cryptography released by the National Institute of Standards and Technology (NIST) [1] indicate that lattice-based cryptographic schemes have significant advantages in both security and efficiency and have become one of the mainstream technical routes [2]. Consequently, students majoring in information security will inevitably encounter the theory of lattice-based cryptography in subsequent courses and research practice, and the formation of this capability depends on the corresponding mathematical foundations.

The existing course Foundations of Information Security Mathematics mainly revolves around number theory and abstract algebra, which effectively supports classical public-key cryptosystems such as RSA, ElGamal, and elliptic curve cryptography. However, the mathematical foundations of lattice-based cryptography are built on vector spaces, and their core problems originate from the SVP, the closest vector problem, and the learning with errors problem. These involve linear representations of vectors, norms and distances, properties of bases, and modular polynomial spaces. Although the current curriculum includes related topics such as polynomial rings, it lacks a systematic presentation of linear space structures and their geometric meanings, making it difficult for students to establish the mathematical cognition required to understand lattice-based cryptography.

From the perspective of talent cultivation, introducing the mathematical foundations of lattice-based cryptography into this course helps extend the frontier of the curriculum [3], enables students to gradually develop knowledge preparation for post-quantum cryptography while learning traditional mathematical content, and lays the necessary foundation for subsequent cryptography courses.

2. Mathematical Foundations Supporting Lattice-Based Cryptography

The security of lattice-based cryptography is based on standard hard problems defined over discrete point sets in vector spaces and their metric properties [4], rather than on number-theoretic or finite algebraic structures. For curriculum design, it is therefore necessary to organize the mathematical foundations around the structure of lattices and the origin of their computational hardness. These foundations can be described at four levels—linear space structures, Euclidean metric structures, algebraic representations of lattices, and lattice-based hard problems—which together provide a framework for understanding the security of post-quantum cryptography.

(1) Linear Space Structures

A lattice can be represented as the set generated by integer linear combinations of several linearly independent vectors. Therefore, the linear representation of vectors constitutes the foundation for understanding lattice structures. At this level, the core teaching content includes the basic concepts of vector spaces, linear independence, bases and dimension, as well as coordinate representations of vectors with respect to different bases.

Unlike traditional linear algebra teaching, which is centered on matrix operations, the emphasis here is on the meaning of representation structures. The same vector has different coordinate representations under different bases, leading to changes in coordinate length and geometric properties. For example, in a two-dimensional space, if two different sets of basis vectors generate the same lattice, it can be observed intuitively that when the basis vectors are nearly orthogonal, the coordinate representation of lattice points is compact, whereas when the basis vectors are nearly linearly dependent, the coordinate length of the same lattice point becomes significantly larger. This phenomenon reveals the relationship between representation and problem difficulty and provides the cognitive basis for understanding lattice basis reduction.

From a structural perspective, a lattice can be viewed as the discrete point set obtained by applying the linear mapping defined by a basis matrix to integer vectors. This viewpoint helps students grasp the generation mechanism of lattices as a whole. On this basis, inner products and vector norms are introduced to provide metric tools for describing lengths and angles rather than for complex computation. The teaching focus of the Gram-Schmidt orthogonalization process is not on algorithmic derivation but on a geometric explanation showing that changing the basis representation can improve the geometric properties of vectors, thereby laying the foundation for understanding the concept of a "good basis".

(2) Metric Structures in Euclidean Space

The criteria for lattice hard problems depend on vector lengths and distances between vectors. Therefore, it is necessary to introduce metric relations in Euclidean space on the basis of linear space structures. The teaching content at this level includes the Euclidean norm, vector distance, and the optimization problems induced by them. In this sense, the shortest vector problem can be interpreted as a norm minimization problem, while the closest vector problem corresponds to a distance minimization problem. Such a formulation unifies the computational hardness in lattices as optimization problems in a metric framework.

In teaching, one may begin with the visualizable case of a two-dimensional Euclidean space and introduce the concept of the closest vector problem through the task of "finding the lattice point closest to a given point", so that students understand the role of distance as the criterion for problem solving. When the dimension is extended from two to higher dimensions, although the generation of the lattice remains known, vector lengths gradually concentrate around a typical value, and most vectors lie in a thin shell at approximately the same distance from the origin. As a result, vectors satisfying the minimum-length or minimum-distance condition become extremely rare, and the corresponding optimization problems quickly become computationally infeasible. Even when relaxed to the problem of finding an approximately shortest vector, the problem remains computationally hard in high dimensions [5].

Through this process, the notion of distance in analytic geometry is transformed into a metric tool for characterizing computational complexity, enabling students to recognize that the security of lattice-based cryptography originates from the metric structure of high-dimensional spaces rather than from the complexity of operational procedures.

(3) Algebraic Structures of Lattices

At the algebraic level, a lattice can be regarded as a point set generated by integer linear combinations of several vectors, and therefore possesses both linear structure and discreteness. In theoretical analysis, the coefficients are taken from the set of integers, whereas in practical cryptographic schemes, in order to ensure computational efficiency and implementability, operations are usually carried out over the finite ring Z_q. In this way, the entire space becomes a finite discrete linear structure.

On this basis, a lattice can be further interpreted as the discrete point set obtained by applying a linear mapping to integer vectors. This formulation not only preserves its linear structure but also gives the generation process a clear algebraic meaning, which helps students grasp the relationship between lattices and vector spaces from a global perspective.

Most mainstream lattice-based cryptographic schemes are built upon structured lattices [6] whose operational spaces are typically given by modular polynomial rings. In this context, a polynomial can be represented by its coefficient vector, and the set of polynomials forms a finite-dimensional discrete linear space modulo q. For example, in the ring $Z_q[x]/(x^n+1)$, a polynomial of degree less than n corresponds one-to-one to an n-dimensional coefficient vector, and polynomial addition corresponds to vector addition. Multiplication by a fixed polynomial is equivalent, in the coefficient representation, to applying a fixed linear transformation to the corresponding vector, and the result remains in the same modular vector space. Therefore, modular polynomial multiplication can be viewed as a linear operation in a finite modular linear space.

From this perspective, the operations in Ring-LWE are essentially carried out in a finite modular discrete linear space, which places the polynomial ring structure and the linear structure of lattices within a unified linear-algebraic framework.

(4) Computationally Hard Problems Based on Lattices

On the basis of the linear representation structure, the metric structure in Euclidean space, and the finite modular algebraic structure described above, several typical computationally hard problems can be introduced as an integrated application background, including the Shortest Vector Problem (SVP), the Closest Vector Problem (CVP), and the Learning With Errors (LWE) problem. The purpose of introducing these problems is not to discuss specific solution algorithms, but to reveal the computational origin of the security of lattice-based cryptography and to enable students to understand how the security of cryptographic schemes is built upon the aforementioned mathematical structures.

The SVP asks for the nonzero lattice vector of minimum length in a given lattice, while the CVP requires finding the lattice vector closest to a given target vector. From the perspective of metric structures, these two problems correspond to norm minimization and distance minimization, respectively. In teaching, one may start from the two-dimensional case and present the solution process through geometric methods. When the dimension is extended to higher-dimensional spaces, the number of vectors satisfying the minimum-length or minimum-distance condition decreases rapidly as the dimension grows. Even when the task is relaxed to finding an approximately shortest vector, these problems remain computationally intractable under existing computational models. Through this contrast between low- and high-dimensional cases, students can recognize that the source of computational hardness lies in the structural change of the space rather than in the increased complexity of computation.

On this basis, the Learning With Errors problem can be introduced as the task of solving a system of noisy linear relations of the form $a_i{\cdot}s+e_i$ (mod q) where a_i are known vectors and the error terms e_i are sampled from a small random distribution, and the goal is to recover the unknown vector s. In the absence of errors, the problem can be solved directly using linear algebraic methods; once the error terms are introduced, the original solution methods no longer apply, making the problem computationally hard under current models. Furthermore, the hardness of LWE can be related to the SVP in q -ary lattices through appropriate constructions [7], thereby establishing a unified connection with the geometric structure at the computational level.

By interpreting SVP and CVP as optimization problems in a metric sense and LWE as a noisy linear structure in a finite modular discrete linear space, a unified framework is established that connects linear representation, metric, and finite modular algebraic structures to cryptographic security. Within this framework, students can recognize that the security of lattice-based cryptography does not arise from the complexity of operational procedures, but from the computational intractability of the associated problems in high-dimensional discrete linear spaces, thereby unifying the computational hardness with the underlying geometric structures.

3. Cognitive Difficulties in Teaching the Mathematical Foundations of Lattice-Based Cryptography

Compared with the number-theoretic and abstract algebraic content in traditional courses on the Foundations of Information Security Mathematics, the mathematical foundations related to lattice-based cryptography exhibit significant differences in both knowledge form and cognitive mode. In teaching practice, these differences are not primarily reflected in an increase in computational complexity, but rather in the cognitive transitions students must undergo when understanding the underlying mathematical structures.

(1) The transition from continuous spaces to discrete point sets constitutes the primary obstacle to understanding lattice structures. In existing courses, vector spaces are usually introduced in the context of continuous spaces. In contrast, a lattice is a discrete point set generated by integer linear combinations, and the object of study is no longer the entire space but the subset determined by a specific generating relation. This shift from a "continuous space" to a "discrete structure" makes it difficult for students to form a global and intuitive understanding of lattices.

(2) The shift from operation-oriented to representation-oriented thinking increases the difficulty of comprehension. Traditional linear algebra teaching focuses on matrix operations and the solution of linear systems, leading students to view linear structures mainly as computational tools. In the mathematical foundations of lattice-based cryptography, however, the basis representation of vectors and the variation of coordinate lengths under different bases become the key source of problem hardness, and the "form of representation" itself becomes

the central object of study. This transition from "how to compute" to "how to represent" challenges students' established cognitive patterns.

(3) The transition from low-dimensional geometric intuition to high-dimensional structures makes it difficult for students to understand the origin of computational hardness. In two- or three-dimensional spaces, the shortest vector problem and the closest vector problem can be understood geometrically and intuitively. In high-dimensional spaces, however, the concentration of vector lengths and the sparsity of short vectors cannot be directly visualized. The absence of geometric intuition easily leads students to attribute the hardness merely to an increase in computational workload, rather than to the intrinsic structure of the space.

(4) The transition from exact linear structures to noisy linear structures is another major obstacle in understanding the Learning With Errors problem. In traditional linear algebra, systems of linear equations are typically solved exactly. In LWE, the introduction of random errors renders the original solution methods ineffective. Students must accept the conclusion that "a linear relation exists but cannot be efficiently solved", which contradicts their prior experience. This cognitive conflict is crucial for understanding the hardness of LWE.

(5) The introduction of finite modular discrete linear spaces changes students' existing understanding of algebraic structures. In abstract algebra courses, polynomial rings mainly appear as objects of algebraic operations. In lattice-based cryptography, however, polynomials are simultaneously viewed as vector representations, and their multiplication corresponds to linear transformations. This shift in the role of the same mathematical object in different contexts makes it difficult for students to establish a unified structural understanding.

In summary, the main difficulty in teaching the mathematical foundations of lattice-based cryptography does not lie in the increase in the number of knowledge points, but in a series of cognitive transitions that students must accomplish, including the shift from continuous to discrete structures, from computation to representation, from low-dimensional intuition to high-dimensional structures, and from exact linear relations to noisy linear structures. These cognitive transitions constitute the key issues that need to be addressed in subsequent instructional design.

4. Instructional Path and Teaching Strategy Design

To address the cognitive transitions required for the mathematical foundations of lattice-based cryptography, the course adopts a progressive introduction of transitional content to guide students from a computation-oriented learning mode to a structure-oriented one.

(1) Cognitive Transition from Continuous Spaces to Discrete Structures

To help students move from continuous vector spaces to discrete lattice point sets, teaching begins with the geometric intuition of lattices in two dimensions. Using a planar point lattice generated by two linearly independent vectors, students observe the difference between the point set produced by integer linear combinations and the entire plane, thereby understanding the structural characteristics of a lattice as a discrete subset. This is then connected to the basis representation of vector spaces, enabling students to recognize that a lattice is not a new mathematical object but the result of imposing integer coefficient constraints on a linear space. The emphasis is not on formal definitions but on developing a global structural intuition, shifting students' perspective from "all vectors in a space" to "a discrete point set determined by generating vectors".

(2) From Operation-Oriented to Representation-Oriented Instruction

Since students tend to view linear algebra primarily as a tool for matrix computation and equation solving, complex calculations are no longer the main training objective. Instead, concrete examples are constructed to show how vector representations affect problem properties. In class, the same two-dimensional lattice is represented with two different bases: one nearly orthogonal and the other nearly linearly dependent. By computing the coordinate representations of the same lattice vector under both bases, students directly observe the difference in coordinate lengths and realize that the form of representation influences problem difficulty.

Students are then guided to consider how changing the basis can make representations more compact, which leads to the introduction of the Gram-Schmidt process. The focus is not on computational procedures but on geometric illustrations of how vector lengths and angles change after orthogonalization. This helps students understand that the essence of the process is to construct a nearly orthogonal basis to improve representations. The idea is subsequently linked to the notion of a "good basis" in lattices, allowing students to see lattice basis reduction as an extension of representation optimization in discrete structures rather than as a new algorithmic topic.

Through this progression from concrete examples to the concept of basis transformation, students shift from the perception that "linear algebra is for computation" to the structural understanding that "representations determine problem properties", which prepares them for understanding the origin of lattice hardness.

(3) From Low-Dimensional Intuition to High-Dimensional Structures

To address the difficulty students encounter in understanding the origin of hardness in high-dimensional spaces, a cognitive path from low-dimensional intuition to high-dimensional structures is adopted. First, concrete lattice examples are constructed in two dimensions, where the distribution of lattice points is visualized and tasks such as finding the shortest nonzero lattice vector and the closest lattice point to a given target are assigned. Students can solve these problems geometrically and develop an intuitive understanding of optimization in a metric sense.

After establishing low-dimensional intuition, graphical representation is replaced by numerical observations. By comparing vector lengths in spaces of increasing dimension through simple numerical experiments, students observe that vector lengths gradually concentrate within a narrow range and that vectors significantly shorter than the average become extremely rare. This data-driven observation allows students to understand that the hardness of SVP and CVP originates from the structure of high-dimensional spaces rather than from the increase in computational steps.

By contrasting the low-dimensional strategy of enumerating candidate vectors with the exponential growth of candidates and the lack of geometric intuition in high dimensions, students recognize why low-dimensional geometric methods fail. Throughout this process, low-dimensional intuition serves as a cognitive anchor, and the sequence "visualized examples –numerical phenomena–structural explanation" helps students interpret lattice hard problems as structural problems in high-dimensional discrete spaces.

(4) From Exact Linear Relations to Noisy Linear Structures

To help students understand the origin of the hardness of the Learning With Errors problem, teaching starts from their familiar experience of solving linear systems. A consistent overdetermined system of linear relations modulo q is first presented, allowing students to recover the unknown vector using elimination or matrix methods and reinforcing the prior belief that exact linear relations can always be solved.

Next, small random perturbations are introduced on the right-hand side while keeping the coefficient matrix unchanged. Students then find that no vector can satisfy all equations simultaneously and that the original solution methods fail. This contrast shows that the hardness is not caused by more unknowns or more complicated procedures, but by the destruction of exact consistency due to noise, transforming the problem from deterministic solving to recovering hidden structure under noise.

It is further explained that the linear relations still exist in an approximate sense and that the task shifts from finding an exact solution to identifying a linear structure masked by random errors. Through this cognitive conflict between solvable and unsolvable cases, students understand that the security of lattice-based cryptography arises from the intractability of recovering noisy linear structures in finite modular discrete spaces. No formal reductions are involved; instead, the structural comparison between exact and noisy cases connects LWE to the hardness of SVP.

(5) Structured Understanding of Finite Modular Discrete Linear Spaces

To address the role shift of polynomial rings between abstract algebra and lattice cryptography, representation-based learning tasks are designed to establish a structural understanding of their operational space. First, modular polynomials of degree less than *n* are introduced, and students are guided to represent them as length-*n* vectors by arranging their coefficients in order. By comparing the results of polynomial addition with those of vector addition, students recognize that the set of polynomials can be viewed as the set of elements of a finite modular vector space. Then, by multiplying several polynomials by a fixed polynomial and observing the linear relations among the coefficients and the preservation of dimension, students understand that modular polynomial multiplication corresponds to a linear transformation in the coefficient representation and remains closed in the same finite modular space.

This representation is then compared with vector-based linear relations and extended to the polynomial form of Ring-LWE, enabling students to see that its operations are another expression of linear operations in a finite modular discrete space. Through the progression "vector representation – polynomial representation – structured lattice space", polynomial rings are integrated into the existing linear-space cognition, leading to a unified structural understanding.

5. Conclusion

Focusing on the mathematical foundations underlying lattice-based cryptography, this paper analyzes the direction of content expansion for the course Foundations of Information Security Mathematics in the context of post-quantum cryptography and constructs a knowledge framework consisting of linear space structures, finite modular discrete spaces, noisy linear relations, and polynomial representations. By clarifying the intrinsic connections among these mathematical objects, the core problems of lattice-based cryptography are unified as structural and computational complexity problems in high-dimensional discrete linear spaces, providing a clear mathematical main line for incorporating related content into the foundational course.

At the instructional level, the traditional emphasis on computational training is shifted to the understanding of representation structures and spatial properties. Through the contrast between low- and high-dimensional cases, the comparison between exact and noisy linear relations, and the consistency between vector and polynomial representations, students are guided to understand the origin of lattice hardness on the basis of their existing linear algebra knowledge. As a result, the key mathematical objects in lattice-based cryptography are no longer presented as new computational techniques but are integrated into a unified framework of linear structural understanding.

Acknowledgements

supported by 'the Fundamental Research Funds for the Central Universities (2023MS136)'

References

[1] Alagic, G., et al. (2022). Status Report on the Third Round of the NIST Post-Quantum Cryptography Standardization Process. NIST IR 8413. https://doi.org/10.6028/NIST.IR.8413

[2] Avanzi, R., Bos, J. W., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J. M., Schwabe, P., Seiler, G., & Stehlé, D. (2021). CRYSTALS-Kyber algorithm specifications and supporting documentation (Version 3.02). https://pq-crystals.org/kyber/data/kyber-specification-round3-20210804.pdf

[3] Buchmann, J. A., Butin, D., Göpfert, F., & Petzoldt, A. (2016). Post-quantum cryptography: State of the art. In P. Ryan, D. Naccache, & J.-J. Quisquater (Eds.), The new codebreakers (Lecture Notes in Computer Science, Vol. 9100, pp. 88–108). Springer. https://doi.org/10.1007/978-3-662-49301-4_6

[4] Peikert, C. (2014). Lattice cryptography for the Internet. In M. Mosca (Ed.), Post-quantum cryptography (Lecture Notes in Computer Science, Vol. 8772, pp. 197–219). Springer. https://doi.org/10.1007/978-3-319-11659-4_12

[5] Aggarwal, D., Dadush, D., Regev, O., & Stephens-Davidowitz, N. (2015). Solving the shortest vector problem in $2n$ time using discrete Gaussian sampling. In Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing (STOC '15) (pp. 733–742). Association for Computing Machinery. https://doi.org/10.1145/2746539.2746606

[6] Albrecht, M. R., Player, R., & Scott, S. (2015). On the concrete hardness of Learning with Errors. Journal of Mathematical Cryptology, 9(3), 169–203. https://doi.org/10.1515/jmc-2015-0016

[7] Micciancio, D., & Regev, O. (2009). Lattice-based cryptography. In D. J. Bernstein, J. Buchmann, & E. Dahmen (Eds.), Post-quantum cryptography (pp. 147–191). Springer. https://doi.org/ 10.1007/978-3-540-88702-7_5