Computer Science and Interdisciplinary Research Journal
Volume 3, Issue 1, 2026. ISSN: 3033-1218, DOI: 10.70862/CSIR.2026.0301-18

Conjugational Stability and Symmetry-Invariant Transformations in Artificial Neural Networks

Kostadin Yotov, Emil Hadzhikolev

University of Plovdiv Paisii Hilendarski, 236, Bulgaria Blvd, 4027 Plovdiv, Bulgaria

yotov@uni-plovdiv.bg, hadjikolev@uni-plovdiv.bg

Abstract – This paper examines the properties of algebraically stabilized linear transformations in artificial neural networks (ANNs) under a change of basis and their compatibility with symmetry-invariant architectures. We investigate the conditions under which a stabilizing transformation of the form $W^k = I + \mu D$, defined over integer matrices, preserves its properties under conjugation, i.e., when passing to an equivalent form $W' = B^{-1}WB$ in a new basis. It is shown that such transformations form a closed class with respect to conjugational equivalence and can be employed in environments with diverse geometric and symmetry requirements. Additionally, we analyze the role of symmetric actions on the input space and define conditions for the invariance of transformations with respect to specific groups of transformations. Both formal algebraic properties and the geometric and semantic interpretation of stability in this context are considered. Different types of stability are presented – algebraic, conjugational, and symmetry-based – along with the relationships between them. The results are relevant to the design of ANN architectures with predictable behavior, robustness under transformations, and compatibility with predefined symmetric structures.

Keywords – algebraic stabilization, linear transformations, conjugational equivalence, symmetry-invariant networks, transformation robustness, symmetries and equivariance.

DOI: https://doi.org/10.70862/CSIR.2026.0301-18

Corresponding author: Emil Hadzhikolev,
University of Plovdiv Paisii Hilendarski
Email: hadjikolev@uni-plovdiv.bg

Received: 12 January 2026.
Revised: 14 February 2026.
Accepted: 15 February 2026.
Published: 16 February 2026.

© 2026 Kostadin Yotov, Emil Hadzhikolev. This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

The article is published with Open Access at http://csirjournal.com/

1. Introduction

At the core of modern artificial neural networks (ANNs) lies the process of transforming input information into a representation that is better suited for solving a specific task – whether classification, regression, segmentation, or decision-making. These transformations are typically implemented through linear operators (weight matrices) and nonlinear functions that operate jointly within layers arranged in deep architectures. Each linear layer, represented by a weight matrix W, aims to transform the input vector x into a new linear space such that relevant features are extracted, amplified, or made more distinguishable for subsequent levels of processing.

In this sense, transformation matrices are not merely parameters but active operators that shape the geometry of the representation. They can implement rotations, projections, scalings, changes of basis, and, in a more general setting, induce group actions or invariant structures on the data. It is precisely through these transformations that an ANN processes and reconfigures the input in order to achieve the desired behavior. The question of how these operators behave and what they actually "do" in the representation space is no less important than the model's final performance.

The deeper or more complex an ANN becomes, the more critical it is for the behavior of individual transformations to be controllable and stable. Without clearly defined constraints on the transformation operators, the ANN may enter unstable regimes – such as exploding gradients, loss of information, poor generalization ability, or instability under input perturbations. Particularly in sensitive architectures such as recurrent, residual, quantized, or symmetry-

invariant ANNs, uncontrolled transformation behavior can lead to severe distortions of the semantic structure of the data.

For this reason, in recent years there has been a clear trend toward stabilizing transformations through a variety of approaches: from constraining spectral norms [1–4], through normalization and regularization techniques [5–8], to geometric and algebraic conditions and explicitly constructed structures [9–11]. The main objective of these methods is to introduce structural robustness, ensuring that operators, when applied repeatedly, do not drift beyond certain behavioral bounds and that essential information is preserved. The introduction of algebraically defined stabilizing conditions – such as exponentiation to the identity form or compatibility with group actions – represents a further step in this direction, as it offers not merely numerical control but a constructive framework for predictability and structural guarantees of the transformations.

In our previous study [12], we introduced the concept of *algebraic stabilization* of linear transformations in ANNs. That study analyzes the conditions under which a given linear transformation matrix $W \in \mathbb{Z}^{n\times n}$, composed of integers and satisfying $det(W) \neq 0$, can be stabilized through exponentiation. The notion of an *algebraically stabilized transformation* is defined as follows: for a given natural number $\mu > 1$, coprime with the determinant of W, there exists an exponent $k \in \mathbb{N}$ and a control matrix $D \in \mathbb{Z}^{n\times n}$ such that

$$W^k = I + \mu D, \tag{1}$$

where I denotes the identity matrix. This equality (or its modular representation) defines a stable algebraic configuration in which repeated application of the transformation yields a result differing from the identity by a controlled additive term modulo μ. This approach has practical relevance for ANNs implemented using integer arithmetic, as well as for quantized or hardware-constrained architectures (FPGA/ASIC).

From a theoretical standpoint, we prove that if $gcd(det(W), \mu) = 1$, then there exists an exponent k such that the reduced matrix $\overline{W} = W \ mod \ \mu$ satisfies

$$W^k = I + \mu D. \tag{2}$$

which shows that equality (1) is indeed fulfilled for some matrix $D \in \mathbb{Z}^{n\times n}$. This result constitutes algebraic stabilization by exponentiation, guaranteeing structural robustness of the transformations without the need for optimization-based training or empirical tuning.

The present paper builds upon the results obtained in [12] and aims to analyze the robustness of algebraic stabilization under a change of basis, as well as the invariant properties of the transformations when considered in the context of conjugation. In this way, a path is outlined toward the development of a more general theory of stabilizing blocks.

2. Change of Basis and Symmetry-Invariant ANNs

In the previous stage of this research [12], we introduced the idea of algebraic stabilization of linear transformations through exponentiation of integer-valued weight matrices. It was shown that, under certain conditions, the transformation W^k exhibits behavior equivalent to the identity modulo μ. This property enables the construction of predictable and controllable transformation blocks.

However, one important question remains open: *what happens if we operate in a different basis?*

2.1. The Problem of Basis Dependence

Let us recall that any linear transformation represented by a matrix W depends on the chosen basis of the space. If we change the basis using an invertible matrix $B \in GL_n(Q)$, then the transformation W in the new basis is represented by the conjugated matrix

$$\mathrm{W}' = B^{-1}WB. \tag{3}$$

This operation does not change the linear nature of the transformation, but it does alter its specific form. If we regard W as a block within an ANN architecture, the following question naturally arises: *will the stabilizing behavior of the transformation be preserved after such a change of basis?*

The answer to this question is important in many practical applications – for example, in computer vision or signal processing – where input data are often not presented in a form suitable for direct processing. For this reason, a preliminary linear transformation is frequently applied:

$$x \mapsto B^{-1}x, \tag{4}$$

whose purpose is to:

- orthogonalize the input components (e.g., via PCA or whitening) [13–15];
- decorrelate and normalize the dimensions [13,16];
- map the input into a “natural” space in which symmetries are more easily captured [14,17–19].

If the layers of an ANN are designed with a stabilizing transformation W, then in order to preserve this stability after the linear preprocessing of the input, one must use the conjugated matrix $\mathrm{W}' = B^{-1}WB$. In this way, algebraic stability is retained under the new input structure.

This issue is also directly relevant in transfer learning and fine-tuning scenarios, where transformations must be adapted across representation spaces. [20–23]. In such cases, a model is pre-trained on one dataset and subsequently applied to another (for example, transfer from English to German in NLP, or from daytime to nighttime images). This often requires adapting coordinate representations and/or applying linear transformations to align the semantic structure of the inputs. This, too, constitutes a form of change of basis. To directly reuse the pre-trained weights W without losing their stabilizing properties, it is necessary to conjugate them into the new space.

In G-equivariant architectures, the ANN output is required to transform in a manner consistent with the symmetries present in the data [24–27]. For example, under a rotation of the input image, the activations should rotate accordingly, while under translation or reflection a covariant response is expected. In this context, many operations are not performed directly in the original basis but rather in a basis adapted to the symmetry group (e.g., via harmonic analysis, Fourier transforms, spherical harmonics, etc.).

A change of basis via B is therefore a mandatory step to enter this group-adapted space. In order for the layers to remain stabilized in this new space, it is important that their form be $\mathrm{W}' = B^{-1}WB$. This guarantees that the stabilization modulo μ, achieved through W, remains valid in the group-transformed basis.

Beyond everything mentioned so far, one must also take into account hardware optimization. In quantized ANNs – especially in FPGA or ASIC implementations – it is often necessary to:

- compress layers via low-rank factorizations [28–31];
- apply structural transformations to achieve better memory layout [32–34];
- work with canonical or special matrix forms (e.g., diagonal, triangular, Toeplitz, etc.) [35–38].

In all of these cases, the original matrix W is not used directly but only after an appropriate linear change of basis. Conjugation $\mathrm{W}' = B^{-1}WB$ enables such implementations without losing previously established properties – including stabilizing ones.

From all the examples presented, it becomes clear that a change of basis via conjugation is not merely a theoretical assumption considered in the theory of ANNs. Rather, it is a natural – and often unavoidable – step in real-world engineering scenarios and in the design of modern ANNs architectures. This implies that if we succeed in combining algebraic stabilization through exponentiation with robustness under basis transformations, we will be able to construct structurally stable layers that:

- can be "transported" across different data representations;
- preserve predictable behavior even under group symmetries and input transformations;
- integrate naturally into computationally efficient hardware implementations.

2.2. Invariance under Conjugation and Transfer of Stability. Analogy with Transfer Learning and ANN Symmetry

From an algebraic point of view, if $W \in \mathbb{Z}^{n\times n}$ satisfies $W^k = I + \mu D$, where $D \in \mathbb{Z}^{n\times n}$, then for any invertible matrix $B \in GL_n(\mathbb{Q})$ we can apply conjugation:

$$\begin{aligned} \mathrm{W}'^k &= (B^{-1}WB)^k = B^{-1}W^kB \\ &= B^{-1}(I + \mu D)B = I + \mu B^{-1}DB. \end{aligned} \tag{5}$$

That is, the new transformation W' also satisfies an equation of the form

$$\mathrm{W}'^k = I + \mu D', \tag{6}$$

where $D' = B^{-1}DB$ and $D' \in \mathbb{Q}^{n\times n}$.

If B has rational (or even integer) entries, then D' will also belong to $\mathbb{Q}^{n\times n}$, or even to $\mathbb{Z}^{n\times n}$ under suitable conditions. This demonstrates that algebraic stabilization is robust under a change of basis via conjugation.

In the context of ANNs, a change of basis can be interpreted as a change in the representation of the input or hidden features. As already mentioned, this idea underlies a wide range of techniques, including:

- transfer learning, where already trained transformations are applied to a new representation of the input;
- group-equivariant neural networks (GNNs), in which the architecture is required to be invariant (or covariant) with respect to groups of transformations – for example, image symmetries such as rotations and translations.

This implies that the proposed algebraic stabilization through exponentiation and conjugation can be naturally employed in such architectures, where the portability of stability across bases (or representations) is essential. The stabilizing behavior is transferable, making it possible to construct ANNs with symmetry-invariant structures in which the properties of one layer propagate to other layers through appropriate transformations.

Figure 1 illustrates precisely such a change of basis. The black circle illustrates the original input of the ANN, while the red ellipse is the result of the transformation by W (the input is transformed into an

ellipse). The blue ellipse (dashed) depicts the result of the transformation by $B^{-1}WB$ – naturally different in the new basis and tilted, yet still an ellipse. Since the corresponding conjugated transformation is applied, the geometric essence of the operation is preserved (it remains an ellipse). This is a powerful tool – both in the presence of symmetries and when adapting to new spaces (transfer learning, equivariant networks, etc.).

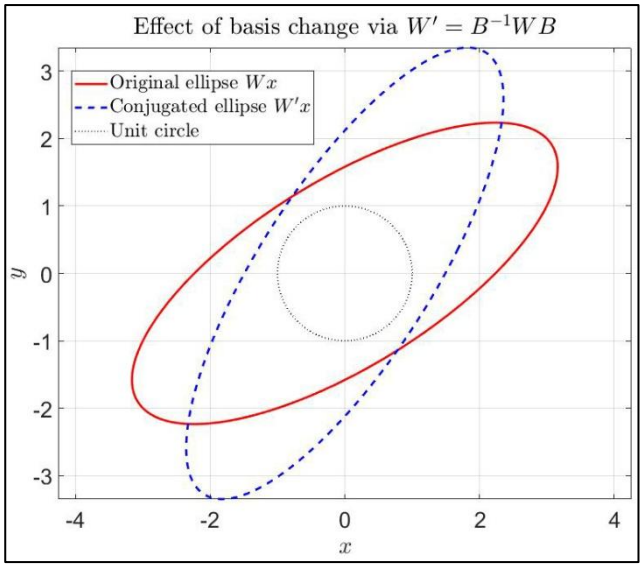


Figure 1. Effect of the change of basis (MATLAB v.2018).

3. Discussion

In the preceding sections of the paper, we examined how the operation of conjugation makes it possible to preserve the stabilizing behavior of linear transformations in ANNs, even when the basis of representation is changed. This property was considered not only out of theoretical interest, but also because of its direct practical relevance – particularly in cases involving data preprocessing, knowledge transfer, symmetry-invariant architectures, and hardware-oriented implementations.

The main advantage (Figure 2) of the proposed concept is that it enables portability of stability: transformations that are stabilized in one basis can be transported via conjugation into another basis without losing their key properties. This is made possible by the fact that the condition for algebraic stabilization is invariant with respect to the conjugation operation. Algebraically, this means that if $W^k = I + \mu D$, then for every square matrix B that is invertible ($detB \neq 0$), we also have:

$$B^{-1}W^kB = I + \mu D', \qquad (7)$$

for an appropriate matrix D'. If B is a nonsingular M-matrix, then $B^{-1} \geq 0$ (elementwise), which ensures monotonicity/positivity of the solutions. In the symmetric case, this coincides with the class of Stieltjes matrices and additionally implies positive definiteness.

Equality (7), expressing preservation of structure under conjugation, enables the construction of structurally stable ANN blocks that can be adapted to different input representations or varying architectural requirements.

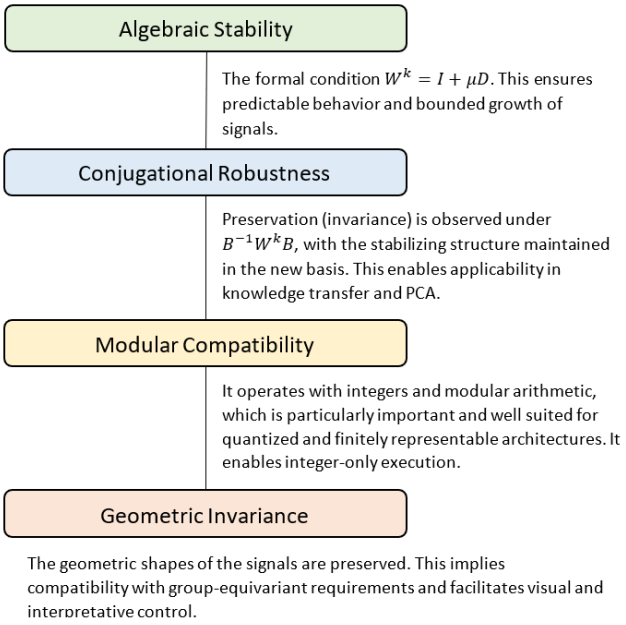


Figure 2. Advantages of algebraic stabilization.

From the perspective of group-equivariant networks, this mechanism enables compatibility with symmetries in the data: if the input data are transformed by rotations, translations, or other linear group actions, the conjugated stabilizing transformations will follow this change and preserve their algebraic properties. This makes it possible to construct symmetry-invariant architectures that are both adaptive and robust.

Moreover, this technique is compatible with real engineering practices. In transfer learning, when trained weights are adapted to new input representations, the possibility of algebraic transfer of stability allows already trained layers to be reused without the need for additional regularization or retraining. A similar situation arises in hardware implementations (e.g., FPGA/ASIC), where the use of special bases – for improved compression, parallelization, or energy efficiency – does not violate the stabilizing properties, provided that the transformations are conjugated correctly.

This type of algebraic stabilization also provides another important tool for visual and interpretative control. It enables model behavior, transformations, or responses (in our case, of an ANN) to be understood, explained, and interpreted in a structured and predictable manner.

In the present context of geometric invariance, this means that when the ANN responds in a particular way to symmetry, rotation, translation, scaling, and similar transformations, these responses are not arbitrary but follow an established regularity that can be interpreted. For example, if we know that the ANN is invariant under rotations, we can explain in advance why rotating the input does not lead to a change in the output. This constitutes interpretative control – knowledge grounded in structure rather than in empirical observation.

The benefits of interpretative control are multiple. It improves model explainability, which is important in sensitive domains such as medicine, autonomous systems, and law. It enables formal verification: when invariance is justified geometrically or algebraically, provable properties of the ANN can be formulated. Finally, it facilitates debugging and error analysis, since we know where deviations in the output can – or cannot – originate.

Of course, the proposed concept also entails certain limitations. Not every change of basis is easy to implement in practice, especially in large-scale ANNs. Moreover, the requirement to maintain a conjugated structure at every layer may introduce additional complexity in architectural design or in the transformation of pre-trained modules. Nevertheless, the advantages clearly outweigh these limitations, particularly for systems that are highly sensitive to the stability of transformations.

Ultimately, through the theoretical framework introduced in [12] and the demonstrated portability of stability across different bases, a foundation is established for constructing ANN structures that are not only effective but also algebraically robust when transitioning across layers or representations. Future research will be directed toward such a construction – namely, a *universal algebraically stable block*.

Finally, we note that the literature frequently employs special classes of Z-matrices in the analysis of invertibility and the sign properties of inverse matrices. Recall that a Z-matrix is a square matrix with nonpositive off-diagonal entries. An M-matrix is a Z-matrix whose eigenvalues have nonnegative real parts (in the nonsingular case, this guarantees invertibility and $B^{-1} \geq 0$ elementwise). An N_0-matrix is an invertible Z-matrix whose inverse is elementwise nonpositive ($B^{-1} \leq 0$) [39]. In future work, we plan to investigate how the choice of matrices B from the classes of Z-, M-, and N_0-matrices influences the construction of algebraically stable blocks and their robustness under finite-precision arithmetic and quantization. Of particular interest is the question of whether these matrix classes are preserved (or how they transform) under conjugation for admissible changes of basis, and whether practically verifiable conditions can be formulated to guarantee preservation of their sign properties. Additionally, we intend to examine the applicability of these matrix classes in the context of monotone/positive dynamical systems, where sign constraints on B^{-1} carry direct physical interpretation.

4. Conclusion

This paper investigates the robustness of algebraically stabilized linear transformations under a change of basis and their compatibility with symmetry-invariant structures in ANNs. The analysis builds upon the theoretical framework introduced in the previous stage of this research, where conditions for algebraic stabilization via exponentiation were defined. Here, we extend this approach by showing that the stabilizing behavior is preserved under conjugation, that is, when the transformation matrix is transferred to a new basis via $\mathrm{W}' = B^{-1}WB$.

We argue that this result is of both theoretical and practical significance. From a theoretical perspective, it demonstrates that the class of algebraically stabilized transformations is closed under conjugation – a structure that plays a central role in group theory, representation theory, and linear algebra. From a practical standpoint, it implies that once a stabilizing matrix has been constructed, it can be adapted to different basis representations without violating its structural properties.

Such a transfer of stability is important in a number of real-world scenarios:

- In transfer learning, where trained parameters are adapted to new input spaces;
- In symmetry-invariant architectures, where transformations must be compatible with groups of symmetries (rotations, translations, etc.);
- In hardware-constrained systems, where changes of basis are used to optimize memory usage, parallelism, or energy efficiency.

The present study also shows that conjugation preserves not only the algebraic form but also the geometric character of the transformations, allowing the semantic structure and symmetric relationships in the input to be maintained. This, in turn, facilitates visual and interpretative control, which is essential for model explainability and reliability – especially in critical domains such as autonomous systems, medicine, or the legal domain.

Despite these advantages, certain limitations should not be overlooked. Changing the basis via conjugation may introduce additional complexity in architectural design, particularly in deep neural networks. Moreover, not every ANN architecture allows for straightforward maintenance of conjugated

transformations across all layers. Nevertheless, the stability and adaptability achieved through this approach make it a highly valuable tool in modern ANN architectural design.

References

[1] Gogianu, F.; Berariu, T.; Rosca, M.C.; Clopath, C.; Busoniu, L.; Pascanu, R. Spectral normalisation for deep reinforcement learning: An optimisation perspective. In *Proceedings of the International Conference on Machine Learning*, 2021; pp. 3734–3744. Retrieved January 8, 2026, from https://proceedings.mlr.press/v139/gogianu21a/gogianu21a.pdf

[2] Keriven, N. Not too little, not too much: A theoretical analysis of graph (over)smoothing. *Advances in Neural Information Processing Systems* 2022, 35, 2268–2281. Retrieved January 8, 2026, from https://proceedings.neurips.cc/paper_files/paper/2022/file/0f956ca6f667c62e0f71511773c86a59-Paper-Conference.pdf

[3] Shi, Z.; Mettes, P.; Maji, S.; Snoek, C.G. On measuring and controlling the spectral bias of the deep image prior. *International Journal of Computer Vision* 2022, 130, 885–908. https://doi.org/10.1007/s11263-021-01572-7

[4] Zhai, S.; Likhomanenko, T.; Littwin, E.; Busbridge, D.; Ramapuram, J.; Zhang, Y.; Susskind, J.M. Stabilizing transformer training by preventing attention entropy collapse. In *Proceedings of the International Conference on Machine Learning*, July 2023; pp. 40770–40803. Retrieved January 8, 2026, from https://proceedings.mlr.press/v202/zhai23a/zhai23a.pdf

[5] Purnomo, A.; Tjandrasa, H. Improved deep learning architecture with batch normalization for EEG signal processing. *JUTI: Jurnal Ilmiah Teknologi Informasi* 2021, 19, 19–27. https://doi.org/10.12962/j24068535.v19i1.a1023

[6] Ziaee, A.; Çano, E. Batch layer normalization: A new normalization layer for CNNs and RNNs. In *Proceedings of the 6th International Conference on Advances in Artificial Intelligence*, October 2022; pp. 40–49. https://doi.org/10.1145/3571560.3571566

[7] Peng, H.; Yu, Y.; Yu, S. Re-thinking the effectiveness of batch normalization and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 2023, 46, 465–478. https://doi.org/10.1109/TPAMI.2023.3319005

[8] Ruhland, J.B.; Masoudian, I.; Heider, D. Enhancing deep neural network training through learnable adaptive normalization. *Knowledge-Based Systems* 2025, 113968. https://doi.org/10.1016/j.knosys.2025.113968

[9] Wang, R.; An, S.; Liu, W.; Li, L. Invertible residual blocks in deep learning networks. *IEEE Transactions on Neural Networks and Learning Systems* 2023, 35, 10167–10173. https://doi.org/10.1109/TNNLS.2023.3238397

[10] Selvaraju, K.; Rajamani, S. New approach to underwater image enhancement using modified residual blocks in generator architecture for improved cycle generative adversarial networks. *Proceedings of the Bulgarian Academy of Sciences* 2024, 77, 73–81. https://doi.org/10.7546/CRABS.2024.01.09

[11] Eskandari, H.; Imani, M.; Moghaddam, M.P. A deep residual network integrating entropy-based wavelet packet ensemble model for short-term electrical load forecasting. *Energy* 2025, 314, 134168. https://doi.org/10.1016/j.energy.2024.134168

[12] Yotov, K., Hadzhikolev, E., & Hadzhikoleva, S. (2026). Algebraic Stabilization of Linear Transformations in Artificial Neural Networks. *Mathematics*, *14*(4), 623. https://doi.org/10.3390/math14040623

[13] Huang, L., Yang, D., Lang, B., & Deng, J. (2018). Decorrelated batch normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 791-800). https://doi.org/10.1109/CVPR.2018.00089

[14] Mobasseri, B. G., & Lulu, A. (2021). Radiometric identification of signals by matched whitening transform. *Sensors*, *21*(24), 8398. https://doi.org/10.3390/s21248398

[15] Zhang, C., Zheng, Y., Guo, B., Li, C., & Liao, N. (2021). SCN: a novel shape classification algorithm based on convolutional neural network. *Symmetry*, *13*(3), 499. https://doi.org/10.3390/sym13030499

[16] Zhang, S., Nezhadarya, E., Fashandi, H., Liu, J., Graham, D., & Shah, M. (2021). Stochastic whitening batch normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 10978-10987). https://doi.org/10.1109/CVPR46437.2021.01083

[17] Xu, J., Kim, H., Rainforth, T., & Teh, Y. (2021). Group equivariant subsampling. *Advances in Neural Information Processing Systems*, *34*, 5934-5946. Retrieved January 8, 2026, from https://proceedings.neurips.cc/paper/2021/file/2ea6241cf767c279cf1e80a790df1885-Paper.pdf

[18] Favoni, M., Ipp, A., Müller, D. I., & Schuh, D. (2022). Lattice gauge equivariant convolutional neural networks. *Physical Review Letters*, *128*(3), 032003. https://doi.org/10.1103/PhysRevLett.128.032003

[19] Mirkes, E. M., Bac, J., Fouché, A., Stasenko, S. V., Zinovyev, A., & Gorban, A. N. (2022). Domain adaptation principal component analysis: base linear method for learning with out-of-distribution data. *Entropy*, *25*(1), 33. https://doi.org/10.3390/e25010033

[20] Shabbir, A., Ali, N., Ahmed, J., Zafar, B., Rasheed, A., Sajid, M., ... & Dar, S. H. (2021). Satellite and scene image classification based on transfer learning and fine tuning of ResNet50. *Mathematical Problems in Engineering*, *2021*(1), 5843816. https://doi.org/10.1155/2021/5843816

[21] Chakraborty, S., Mondal, R., Singh, P. K., Sarkar, R., & Bhattacharjee, D. (2021). Transfer learning with fine tuning for human action recognition from still images. *Multimedia Tools and Applications*, *80*(13), 20547-20578. https://doi.org/10.1007/s11042-021-10753-y

[22] Araújo, R. L., Araújo, F. H. D., & Silva, R. R. E. (2022). Automatic segmentation of melanoma skin cancer using transfer learning and fine-tuning. *Multimedia Systems*, *28*(4), 1239-1250. https://doi.org/10.1007/s00530-021-00840-3

[23] Prottasha, N. J., Sami, A. A., Kowsher, M., Murad, S. A., Bairagi, A. K., Masud, M., & Baz, M. (2022). Transfer learning for sentiment analysis using BERT based supervised fine-tuning. *Sensors*, *22*(11), 4157. https://doi.org/10.3390/s22114157

[24] Bökman, G., Kahl, F., & Flinth, A. (2022). Zz-net: A universal rotation equivariant architecture for 2d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 10976-10985). https://doi.org/10.1109/CVPR52688.2022.01070

[25] Sanborn, S., & Miolane, N. (2023). A general framework for robust g-invariance in g-equivariant networks. *Advances in Neural Information Processing Systems*, *36*, 67103-67124. Retrieved January 8, 2026, from https://proceedings.neurips.cc/paper_files/paper/2023/file/d42523d621194ba54dda098669645f91-Paper-Conference.pdf

[26] Navon, A., Shamsian, A., Achituve, I., Fetaya, E., Chechik, G., & Maron, H. (2023, July). Equivariant architectures for learning in deep weight spaces. In *International Conference on Machine Learning* (pp. 25790-25816). PMLR. Retrieved January 8, 2026, from https://proceedings.mlr.press/v202/navon23a/navon23a.pdf

[27] Kim, J., Nguyen, D., Suleymanzade, A., An, H., & Hong, S. (2023). Learning probabilistic symmetrization for architecture agnostic equivariance. *Advances in Neural Information Processing Systems*, *36*, 18582-18612. Retrieved January 8, 2026, from https://proceedings.neurips.cc/paper_files/paper/2023/file/3b5c7c9c5c7bd77eb73d0baec7a07165-Paper-Conference.pdf

[28] Jaderberg, M., Vedaldi, A., & Zisserman, A. (2014). Speeding up Convolutional Neural Networks with Low Rank Expansions. *British Machine Vision Conference*. BMVA Press. http://dx.doi.org/10.5244/C.28.88

[29] Deng, L., Li, G., Han, S., Shi, L., & Xie, Y. (2020). Model compression and hardware acceleration for neural networks: A comprehensive survey. *Proceedings of the IEEE*, *108*(4), 485-532. https://doi.org/10.1109/JPROC.2020.2976475

[30] Wu, R., Guo, X., Du, J., & Li, J. (2021). Accelerating neural network inference on FPGA-based platforms—A survey. *Electronics*, *10*(9), 1025. https://doi.org/10.3390/electronics10091025

[31] Sui, X., Lv, Q., Bai, Y., Zhu, B., Zhi, L., Yang, Y., & Tan, Z. (2022). A hardware-friendly low-bit power-of-two quantization method for CNNs and its FPGA implementation. *Sensors*, *22*(17), 6618. https://doi.org/10.3390/s22176618

[32] Wang, C., & Luo, Z. (2022). A review of the optimal design of neural networks based on FPGA. *Applied Sciences*, *12*(21), 10771. https://doi.org/10.3390/app122110771

[33] Tasci, M., Istanbullu, A., Tumen, V., & Kosunalp, S. (2025). FPGA-QNN: quantized neural network hardware acceleration on FPGAs. *Applied Sciences*, *15*(2), 688. https://doi.org/10.3390/app15020688

[34] Wang, Y., & Zhao, H. (2025). An Improved Strategy for Data Layout in Convolution Operations on FPGA-Based Multi-Memory Accelerators. *Electronics*, *14*(11), 2127. https://doi.org/10.3390/electronics14112127

[35] Liao, S., Samiee, A., Deng, C., Bai, Y., & Yuan, B. (2019, May). Compressing deep neural networks using toeplitz matrix: Algorithm design and fpga implementation. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 1443-1447). IEEE. https://doi.org/10.1109/ICASSP.2019.8683556

[36] Wang, E., Davis, J. J., Zhao, R., Ng, H. C., Niu, X., Luk, W., ... & Constantinides, G. A. (2019). Deep neural network approximation for custom hardware: Where we've been, where we're going. *ACM Computing Surveys (CSUR)*, *52*(2), 1-39. https://doi.org/10.1145/3309551

[37] Zacchigna, F. G., Lew, S., & Lutenberg, A. (2024). Flexible quantization for efficient convolutional neural networks. *Electronics*, *13*(10), 1923. https://doi.org/10.3390/electronics13101923

[38] Zhang, L., Krestinskaya, O., Fouda, M. E., Eltawil, A. M., & Salama, K. N. (2025). Quantized convolutional neural networks: a hardware perspective. *Frontiers in Electronics*, *6*, 1469802. https://doi.org/10.3389/felec.2025.1469802

[39] Mondal, S., & Sivakumar, K. C. (2025). A Survey of Z-Matrices and New Results on the Subclass of F0-Matrices. In *Inverse Problems, Regularization Methods and Related Topics: A Volume in Honour of Thamban Nair* (pp. 303-327). Singapore: Springer Nature Singapore. https://doi.org/10.1007/978-981-97-7989-5_14