UNIVERSITÀ DEGLI STUDI DI PADOVA

Dipartimento di Fisica e Astronomia "Galileo Galilei"

Master Degree in Physics

Final Dissertation

Consistent gaugings of maximal supergravity in 4d

Thesis supervisor
Prof. Gianguido Dall'Agata

Candidate
Erica Di Bari

Academic Year 2025/2026

Abstract

Maximal supergravity in four dimensions provides a highly constrained and mathematically rigorous environment in which to investigate possible gaugings and structures of vacua. This thesis analyses gaugings of the theory, relying on the embedding tensor formalism, with particular focus on its rank. What makes the embedding tensor unique is that it brings within all information regarding the gauging. Furthermore, it is characterized by strict algebraic consistency dictated by linear and quadratic constraints, that originate from supersymmetry and locality, and their explicit form depends on the $E_{7(7)}$ duality group structure and the symplectic geometry of its fundamental representation.

The study is conducted in the canonical electric frame of the theory. The first part focuses on analytical boundaries, specifically investigating the existence of valid gaugings with an embedding tensor of very low rank. For the rank-1 case, we analytically demonstrate that non-trivial solutions are impossible to obtain. To tackle the significantly more complex higher-rank scenarios, we developed dedicated computational pipelines using Wolfram Mathematica and Python. Utilizing known $E_{7(7)}$ group data, the code systematically tests rank-n cases to identify the lowest rank capable of yielding non-trivial solutions. A frame-independent computational check provides robust evidence that the result is general.

The strategy implemented combines exact mathematical formulations of the constraints with numerical and symbolic verification of the closure conditions, subsequently classifying any valid solutions, comparing them to known ones. While certain heuristic steps within the code depend on precision thresholds, the goal is not to be exhaustive but to explore different computationally adaptable frameworks for inspecting maximal supergravity gaugings, in particular from the perspective of rank stratification. This automated classification serves as a valuable baseline for future improvements in gauging analysis, as well as being intrinsically connected to studies of the vacuum structure.

Contents

Chapter 1

Introduction

The Standard Model of particle physics and General relativity are the pillars on which contemporary physics stands. Despite the fact that both theories have been vastly supported by experimental validation, from the discovery of the Higgs boson to the detection of gravitational waves, they remain separate theories. The absence of a unifying theory is one of the most fascinating challenges that theoretical physics is set on solving. The main concern when handling a quantized theory of General Relativity is its catastrophic behaviour in the ultraviolet. The fact that the coupling constant of the theory, the notorious Newton's constant $G_N \equiv \frac{\hbar c}{M_P^2}$ (where M_P is the Planck mass), possesses a negative mass dimension prevents standard renormalizability. At the Planck scale, the theory fails miserably and loses predictivity. This seems to indicate that General Relativity is an effective field theory valid at low energies only, and whose UV-completion has to be studied cautiously. This has led to the development of deeper, more complex frameworks that aim to solve this incompatibility. The Standard Model itself, however, is not exempt from profound issues either: naturalness problems, such as the value of the neutrino masses and the hierarchy problem, signal that there are phenomena for which the current framework cannot provide a satisfying answer, and one has to go beyond the Standard Model. A fascinating approach to this problem is given by String Theory, which seems to be the leading theoretical model for a UV-description of Quantum Gravity [1]. The replacement of the traditional point-like particle of QFT with strings, one dimensional objects spanning worldsheets in spacetime, make the high-energy description of the theory radically different. UV divergencies are then manageable and an integration of gravity with the other fundamental forces seems feasible. One can look at supergravity as the bridge between the high-energy framework of string theories and the low-energy macroscopic world. Specifically, it is an ordinary Quantum Field Theory that contains both gravity and supersymmetry, the two necessary ingredients to constitute effective string theories. One interesting aspect is that supersymmetry requires the presence of a superpartner for the graviton, called gravitino, a spin-$\frac{3}{2}$ particle. This particle enters in loops that otherwise would cause bad UV divergence, causing bosonic-fermionic group cancellations.

A characteristic of string theories is that they live in higher dimensions: ten dimensions are required to maintain quantum consistency and anomaly cancellations. The universe we live in, however, has only four dimensions. To see how this theory reflects on our 4d reality, therefore, it is necessary to "get rid" of the extra dimensions. This happens through the process of **compactification**, where these unwanted dimensions get "rolled up" into internal manifolds. However, compactification usually leaves behind several massless scalar fields, called moduli, parametrizing the geometrical structure of the compactified manifold. Since they are massless, they should also be mediators of long-range forces that, however, contradict empirical findings. This is the moduli stabilization problem, and it is fixed by using flux compactifications. Without delving too much on the details, these fluxes are basically higher-dimensional, generalized electromagnetic fields. Their action basically consists in wrapping the background quantum field energy around the internal geometry to stabilize the moduli: by integrating over the extra dimensions, this process reflects on the correspondent effective four-dimensional theory, giving it a scalar potential and fixing the vacuum expectation value of the moduli. As a consequence, they are no longer massless, and we have got rid of these unwanted long-range forces.

Supergravity emerges exactly in this way, as a four-dimensional, low-energy product of String Theory. Here, flux compactification gives a non-trivial potential whose v.e.v. that determines the cosmological constant Λ. This shows how a seemingly abstract theory such as String Theory is heavily connected to a crucial observable: by determining the sign of Λ, we know if the theory describes a de Sitter ($\Lambda > 0$), Anti-de Sitter ($\Lambda < 0$) or Minkowski ($\Lambda = 0$) spacetime. However, supergravity theories can be studied on their own as a highly constrained quantum field theory, therefore employing a bottom-up perspective. In particular, $\mathcal{N} = 8$ maximal supergravity is a rich and unique case to probe. It was originally formulated by Cremmer and Julia in 1978 [2] and it became the ground for extensive successive studies. $\mathcal{N} = 8$ supergravity possesses the maximum allowable number of supersymmetries, and thus it is unique in that it is the most symmetric field theory containing gravity that does not involve particle with spins greater than 2. Theories with a higher number of supercharges would inevitably introduce particles with higher spin (such as $\frac{5}{2}$), which are forbidden from having non-trivial interactions by fundamental no-go theorems (most notoriously by the Weinberg-Witten theorem first presented in 1980 [3]). Specifically, they prove that no relativistic quantum field theory possessing a conserved and Lorentz-covariant stress-energy tensor $T_{\mu\nu}$ can accommodate interacting massless particles with spin $s > 1$. Furthermore, while massless particles with spins higher than 2 may theoretically exist, their couplings inevitably vanish in the low-energy limit, preventing them from mediating long-range interactions. Both the graviton and the gravitino, however, remain unaffected, since there is no gauge-invariant and Lorentz-covariant stress-energy tensor for these fields. Of course, this high level of symmetry makes the theory extremely rigid, which probably means that it will not correspond to a realistic model that describes our universe. Nevertheless, for theories such as $N = 4$ SUGRA, the gravity multiplet can in principle couple to infinite matter vector multiplets. While this is a freer theory, the clear advantage of maximal $\mathcal{N} = 8$ SUGRA is that its particle spectrum and geometry are entirely determined and finite.

The maximal theory is characterized by a global symmetry group called $E_{7(7)}$, which acts as a generalized electric/magnetic duality. While the full mathematical group of electric/magnetic swaps for the vectors is the larger symplectic group $Sp(56, \mathbb{R})$, $E_{7(7)}$ is embedded inside it. This distinction is critical: while the larger $Sp(56, \mathbb{R})$ group is a symmetry only of the equations of motion and the Bianchi identities, $E_{7(7)}$ is a true symmetry of the entire theory. Furthermore, the scalar fields in the theory form a precise structure: the coset space $\mathrm{E}_{7(7)}/\mathrm{SU}(8)$. In its original, ungauged formulation, $\mathrm{E}_{7(7)}$ is purely global, and one characteristic is that no scalar potential appears at all. When implementing a gauging procedure, a subgroup $\mathrm{G}_0 \subset \mathrm{E}_{7(7)}$ is promoted to a local gauge symmetry, and this leads to the presence of a potential. However, handling the gauging procedure as one typically does (in the Standard Model, for example, where one chooses a gauge subgroup a priori) is not convenient in this case. Instead of picking subgroups by hand and breaking covariance for the start, it would be easier to encode all gauging steps and calculations on a single tool that defers the particular gauge choice to the end. This is why the **embedding tensor formalism** was constructed and applied for the first time to maximal supergravity by Samtleben, de Wit and Trigiante in [4]. The embedding tensor is an object ${\Theta_M}^\alpha$ which carries both the adjoint **133** representation indices (α) and the fundamental **56** (M) of $\mathrm{E}_{7(7)}$. It is related to the gauge generators X_M via:

$$X_M = {\Theta_M}^\alpha t_\alpha \tag{1.1}$$

where t_α are the generators of $\mathrm{E}_{7(7)}$ in the adjoint representation. The embedding tensor will then enter the covariant derivative $D_\mu = \partial_\mu - A_\mu^M {\Theta_M}^\alpha t_\alpha$ Furthermore, to ensure that the gauging procedure maintains consistency, the embedding tensor has to satisfy three (not independent) conditions [5]: a linear constraint, for supersymmetry preservation, and the quadratic closure constraint, which verifies that the generators properly form a closed algebra, and finally, the locality constraint. This last constraint arises from the fact that we are dealing with **duality**. In the un-gauged theory, the 56 vector fields are arranged in 28 electric and 28 magnetic duals, and this is governed by a global $Sp(56, \mathbb{R})$ symplectic symmetry. Due to this underlying structure, it is impossible to construct a Lagrangian that simultaneously contains all 56 vectors, as it would incorrectly double the degrees of freedom. This condition also implies that, by performing an $\mathrm{E}_{7(7)}$ transformation, one can always find an electric frame where all magnetic components of the embedding tensor vanish.

Basically, all gauging information is encoded in the embedding tensor: every deformation of the theory is stored inside it. This reflects also on the potential, which emerges from gauging as a function of the embedding tensor. Therefore, once the latter is specified, one can analyze the structure of vacua. This is exactly why this thesis focuses on the study of the embedding tensor: understanding its properties and realizations is a fundamental step in gaining deeper understanding of the vacua structure of the maximal theory. These conclusions can then also be subject of study in the context of flux compactifications and string theory, as explained before.

Here we focus on a specific analysis for the embedding tensor: rank-by-rank stratification. The main question is whether it is possible to find the smallest rank that admits any gauging at all. An analytical approach is tried first, but the case becomes quickly complicated and requires computational methods. Therefore, different pipelines are implemented to conduct this scan. First, a proof for the $\mathrm{SL}(8, \mathbb{R})$ electric frame is provided, but to prove frame-invariance across $\mathrm{E}_{7(7)}$, a study of the minimal allowed support of the linear constraint is added. After this analysis, the pipelines continued the scan for higher ranks in the $\mathrm{SL}(8, \mathbb{R})$ frame, to understand the distribution of gaugings and also to explore possible computational methods that might serve to find gaugings starting from an embedding tensor of a certain rank. The check for validity is made by comparing the results of these scans with the ones in [6], which catalogued the whole family of $\mathrm{CSO}(\mathrm{p}, \mathrm{q}, \mathrm{r})$ gaugings in $\mathrm{SL}(8, \mathbb{R})$. The first scan uses a stochastic approach for the calculation of the closure condition, and has severe limitations for gaugings that lie in very narrow basins; this suggested the implementation for another computational pipeline, the exact scan, which tries to retrieve entire families of gaugings instead of single-point solutions. This scan proved more useful but still limited since it has some heuristic steps, and its complexity varies with the symmetry characteristics of the embedding tensor in the considered frame. These scans are tentative approaches and the conclusions reached offer stimuli for possible further improvements and studies, especially for the exact scan. Finally, a brief systematic analysis of the mass basis case was executed, employing the Lie algebra contraction formalism developed in [7, 8]. Contractions provide a rigorous framework for taking singular limits of the underlying exceptional symmetry groups. A detailed rank analysis of these contracted algebras allows us to classify the structural transitions of the system, tracking the specific mass basis configurations.

Concretely, in Chapter 2 we present the general framework required for gauging supergravity theories. The process of promoting duality groups to local gauge symmetries is discussed in its fundamental steps, without involving maximal supergravity from the start. In this context, the embedding tensor formalism is constructed, as well as its emerging constraints. The scalar manifold is studied in a general supergravity context, together with the scalar potential and its ties to the embedding tensor. A final presentation of the purely electric frame convention concludes the chapter.

In Chapter 3, this general framework is specialized to $\mathcal{N} = 8$ maximal supergravity, starting with its field content and ungauged lagrangian, continuing with the discussion of the gauging process with all its nuances, such as consistency with supersymmetry. The embedding tensor and the scalar potential are also finally presented in the context of the maximal theory.

After establishing the theoretical background, Chapter 4 describes the methods used to obtain the results for this thesis, and in particular focuses on how the no-go theorem for the minimum admissible rank of the embedding tensor was formulated. After showing examples of analytical approaches for the lowest ranks, including a demonstration in the canonical electric frame for the rank-1 case, the pipelines for the two scans (stochastic and exact) implemented are introduced. The code conducts the analysis in the $\mathrm{SL}(8, \mathbb{R})$ known frame, which provides the chance to compare results with the study conducted in [6], serving as a guideline for a consistency check for the lowest-rank gauging found and to understand how much these pipelines actually prove useful for gauging classifications, beyond the initial exact check for the minimum rank admissible.

In Chapter 5, finally, all results are reported, including the statement for the no-go theorem. A robust generalization to this statement, which is proved exactly only within the $\mathrm{SL}(8, \mathbb{R})$ frame, is found by analyzing the minimum support allowed for the linear constraint. Despite suffering from computational limitations, it is still a strong corroboration of our finding that extends this result across $E_{7(7)}$, since it is devoid of a choice of frame. The mass basis case is finally presented and analyzed.

Chapter 2

Gauging and the embedding tensor

The foundation of any gauged supergravity consists in the promotion of a subgroup of the theory's global symmetries to local (gauge) symmetries, through the introduction of gauge fields and their minimal coupling. This is implemented via the covariant derivative, which encodes how matter fields transform under the gauged symmetry group. However, in maximal supergravity, gauging goes beyond these standard steps: to preserve supersymmetry, localizing the gauge group triggers non-trivial deformations of the theory. This becomes manifest through the presence of non-canonical kinetic terms and non-minimal field couplings. A fundamental consequence of gauging in maximal supergravity is also the generation of a scalar potential, otherwise absent in the theory. In this chapter I will present the general process of gauging a theory, moving later onto the motivation behind the employment of the embedding tensor formalism for the gauging of maximal supergravity, and in the context of this thesis specifically: the highly constrained setting induced by it will, in fact, provide a very useful tool to work with this theory in dimension 4. The main references used throughout this chapter are [4], [9] and [5].

2.1 Gauging global symmetries

Before introducing the framework of the embedding tensor, it is instructive to present the gauging process. Considering the case of a generic field theory with global symmetry group G, *gauging* implies the selection of a subgroup $G_0 \subset G$ and its promotion to a local symmetry.

The first step is to consider a Lagrangian $\mathcal{L}[\Phi]$ that depends on a set of fields $\Phi = \{A_\mu^M, \phi^i, \psi\}$ (vectors, scalars, and fermions respectively) that is invariant under the action of a Lie group G,

$$\delta_\epsilon \Phi = \epsilon^\alpha \left(t_\alpha \Phi\right), \qquad \epsilon^\alpha \in \mathbb{R}, \tag{2.1}$$

where t_α are the generators of $\mathfrak{g} = \mathrm{Lie}(G)$ satisfying

$$[t_\alpha, t_\beta] = f_{\alpha\beta}{}^\gamma \, t_\gamma, \tag{2.2}$$

with $f_{\alpha\beta}{}^\gamma$ the structure constants of G. By the well-known Noether's theorem, each independent continuous symmetry gives rise to a conserved current j_μ^α and an associated conserved charge. The representation $\mathcal{R}_v$ under which the vector fields A_μ^M transform is particularly relevant, so that under

G

$$\delta_\epsilon A_\mu^M = \epsilon^\alpha \, (t_\alpha)^M{}_N \, A_\mu^N \,, \tag{2.3}$$

where $(t_\alpha)^M{}_N$ is the matrix representative of t_α in $\mathcal{R}_v$. The idea behind gauging is to allow the symmetry parameters to depend on a spacetime point, $\epsilon^\alpha \to \epsilon^\alpha(x)$. Under a local variation the kinetic terms of the fields are no longer invariant, because

$$\delta_{\epsilon(x)}\partial_\mu \Phi = \epsilon^\alpha(x)\, \partial_\mu(t_\alpha \Phi) + (\partial_\mu \epsilon^\alpha(x))\, t_\alpha \Phi \,, \tag{2.4}$$

and the second term spoils invariance. The standard remedy is to replace the ordinary derivative with a *gauge-covariant derivative*,

$$D_\mu \Phi = \partial_\mu \Phi - g\, A_\mu^\alpha \, t_\alpha \Phi \,, \tag{2.5}$$

where g is the gauge coupling and A_μ^α are gauge connection one-forms transforming in the adjoint representation of G_0, while the choice of the minus sign is a convention. This formulation requires the representation of G_0 to be faithful, that is, there must be a one-to-one correspondence between the elements of the Lie algebra and their matrix representations acting on the fields. In this way, every local gauge transformation corresponds to a unique shift in the fields, avoiding blind mapping. Invariance of the covariant derivative under local transformations is then restored, provided that the gauge fields transform as

$$\delta_\epsilon A_\mu^\alpha = \partial_\mu \epsilon^\alpha + g\, f_{\beta\gamma}{}^\alpha \, \epsilon^\beta A_\mu^\gamma \,. \tag{2.6}$$

The kinetic term for the gauge fields themselves is given by the Yang–Mills Lagrangian,

$$\mathcal{L}_{\mathrm{YM}} = -\tfrac{1}{4}\, F_{\mu\nu}^\alpha F^{\mu\nu}{}_\alpha \,, \qquad F_{\mu\nu}^\alpha = \partial_\mu A_\nu^\alpha - \partial_\nu A_\mu^\alpha + g\, f_{\beta\gamma}{}^\alpha \, A_\mu^\beta A_\nu^\gamma \,, \tag{2.7}$$

which is invariant under the local gauge symmetry by construction. This is the standard procedure employed for gauging the Standard Model: starting from a theory with global symmetry G, one introduces one gauge boson for each generator of G_0.

Gauging and duality

The situation becomes more subtle in theories where the equations of motion and the Bianchi identities (but not necessarily the Lagrangian) are invariant under a large *duality group*. In the context of supergravity, the global symmetry group G is identified with the **U-duality group** G_U, so in what follows I will refer to G_U. To see where this case differs from standard gauging, consider n_v abelian vector fields with field strengths $F_{\mu\nu}^I$ $(I = 1, \ldots, n_v)$. The Bianchi identities and equations of motion take the unified form

$$\partial^\mu \widetilde{F}_{\mu\nu}^I = 0 \quad \text{(Bianchi)}, \qquad \partial^\mu G_{I\,\mu\nu} = 0 \quad \text{(EOM)}, \tag{2.8}$$

where $\widetilde{F}_{\mu\nu}^I = \frac{1}{2}\varepsilon_{\mu\nu\rho\sigma}F^{I\,\rho\sigma}$ is the Hodge dual and $G_{I\,\mu\nu} \equiv -2\,\partial\mathcal{L}/\partial F^{I\,\mu\nu}$ is the dual field strength. This is the natural generalisation of the familiar case of Maxwell theory: in vacuum, the equations for $\vec{E}$ and $\vec{B}$ are symmetric under the exchange $\vec{E} \to \vec{B}$, $\vec{B} \to -\vec{E}$, meaning that electric and magnetic fields can be swapped without changing the physics. In the covariant language, this is the rotation $(F_{\mu\nu}, \widetilde{F}_{\mu\nu}) \to (\widetilde{F}_{\mu\nu}, -F_{\mu\nu})$, a discrete $\mathbb{Z}_4$ transformation inside the continuous U(1) of rotations of the pair $(F, \widetilde{F})$. In supergravity [5], $F^I{}_{\mu\nu}$ and $G_{I\mu\nu}$ have the role of generalized electric and magnetic field strengths, where $G_{I\mu\nu}$ reduces to $F^I{}_{\mu\nu}$ in the free field limit, but is dressed non-linearly by the scalar-dependent matrix $\mathcal{N}_{IJ}(\phi)$, presented later in the section. To determine what group leaves this structure invariant, consider first that the electric field strengths and their magnetic duals can be paired into a single $2n_v$-component column vector in the fundamental representation of the duality group:

$$\mathcal{F}^M_{\mu\nu} = \begin{pmatrix} F^I_{\mu\nu} \\ G_{I\mu\nu} \end{pmatrix} \tag{2.9}$$

The unified field equations now look like:

$$d\mathcal{F}^M = 0 \implies \begin{pmatrix} dF^I \\ dG_I \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \tag{2.10}$$

where d denotes the exterior derivative. Let us consider a linear duality transformation that acts on $\mathcal{F}$ as:

$$\mathcal{F}'^M_{\mu\nu} = \mathcal{S}^M{}_N \mathcal{F}^N_{\mu\nu} \tag{2.11}$$

where $\mathcal{S}$ is a $2n_v \times 2n_v$ real matrix. In order to preserve the structural mapping between the electric field strengths and their magnetic duals, when rotating the EOM and Bianchi identities into each other, any physically valid transformation matrix $\mathcal{S}$ must leave the constant, antisymmetric symplectic matrix invariant:

$$\Omega = \begin{pmatrix} 0 & \mathbb{I}_{n_v} \\ -\mathbb{I}_{n_v} & 0 \end{pmatrix} \tag{2.12}$$

And the invariance for the system of equations is guaranteed under $S \in GL(2n_v, \mathbb{R})$. However, since the $G_{I\mu\nu} \equiv -2\frac{\partial \mathcal{L}}{\partial F^{I\mu\nu}}$ field strength depend on the Lagrangian, and this relation must be preserved as well under S-transformations, additional constraints on this matrix emerge. Looking at the components for S, the emerging constraints are:

$$S = \begin{pmatrix} A & B \\ C & D \end{pmatrix}, \tag{2.13}$$

$$C = C^T, \quad B = B^T, \quad A = -D^T. \tag{2.14}$$

These invariance constraints are equivalent to the conservation of the so-called **symplectic form**:

$$\Omega_{MN} = \begin{pmatrix} 0 & \mathbf{1}_{n_v} \\ -\mathbf{1}_{n_v} & 0 \end{pmatrix}, \qquad \mathcal{S}^T\,\Omega\,\mathcal{S} = \Omega. \tag{2.15}$$

Its inverse Ω^{MN} (defined by $\Omega^{MN}\Omega_{NP} = -\delta^M{}_P$) is used to raise symplectic indices. Therefore, the general linear group is restricted to the symplectic group $\mathrm{Sp}(2n_v, \mathbb{R})$. The relation between G_I and F^J is not free, but as mentioned before, it is governed by $\mathcal{N}_{IJ}(\phi)$, a complex symmetric matrix:

$$G_{I\,\mu\nu} = \mathrm{Re}\,\mathcal{N}_{IJ}\,F^J_{\mu\nu} + \mathrm{Im}\,\mathcal{N}_{IJ}\,\widetilde{F}^J_{\mu\nu}. \tag{2.16}$$

$\mathcal{N}_{IJ}$ is a complex symmetric matrix whose imaginary part $\mathrm{Im}\,\mathcal{N}$ is negative definite, which ensures the vector kinetic term has the correct (positive) sign.

While the combined Bianchi–EOM system is $\mathrm{Sp}(2n_v, \mathbb{R})$-invariant, this is not generally true for the Lagrangian. An infinitesimal $\mathrm{Sp}(2n_v, \mathbb{R})$ transformation, in fact, shifts it by

$$\delta\mathcal{L} = \tfrac{1}{4} C_{IJ} F^I_{\mu\nu} \widetilde{F}^{J\,\mu\nu} + \tfrac{1}{4} B^{IJ} G_{I\,\mu\nu} G_J{}^{\mu\nu} + \delta\mathcal{L}_{\text{rest}}, \tag{2.17}$$

where B^{IJ} and C_{IJ} are the off-diagonal blocks of the infinitesimal symplectic transformation generator (for simplicity, the letters are the same of the global full transformation (2.13). Therefore, only a proper subgroup of $\mathrm{Sp}(2n_v, \mathbb{R})$, called **U-duality group** G_U preserves invariance: it consists of the set of transformations that can be compensated by a simultaneous action on the scalar and fermion fields, managing to leave the complete Lagrangian invariant, so that $\delta\mathcal{L}_{rest} = 0$. These are precisely the isometries of the scalar manifold $\mathcal{M}_{\text{scalar}}$,

$$\delta_A \phi^i = \xi^i_A(\phi), \qquad G_U \cong \mathrm{Iso}(\mathcal{M}_{\text{scalar}}), \tag{2.18}$$

where $\xi^i_A(\phi)$ are the Killing vector fields on $\mathcal{M}_{\text{scalar}}$, the vector fields that generate infinitesimal isometries. The index A labels the generators of G_U and i labels the scalar field directions. Here, $G_{\text{global}} = G_U \times G_{\text{inert}}$, a direct product of two commuting factors. Here G_{inert} denotes the subgroup of the global symmetry that acts on fields which are neutral under the duality group, with no associated vector field and that do not participate in any electric–magnetic duality rotation. To see why G_U coincides with the isometry group of the scalar manifold, recall that the vector kinetic Lagrangian takes the form

$$\mathcal{L}_{\text{vec}} \sim \mathrm{Im}\,\mathcal{N}_{IJ}(\phi)\, F^I_{\mu\nu} F^{J\,\mu\nu} + \mathrm{Re}\,\mathcal{N}_{IJ}(\phi)\, F^I_{\mu\nu} \widetilde{F}^{J\,\mu\nu}, \tag{2.19}$$

Where $\mathcal{N}_{IJ}$ is the kinetic matrix in (2.16). This relation can be rewritten by introducing the self-dual and anti self-dual field strengths, $F^{\pm J}_{\mu\nu} \equiv \frac{1}{2}\left(F^J_{\mu\nu} \mp i\tilde{F}^J_{\mu\nu}\right)$:

$$G_I^{+\mu\nu} = \mathcal{N}_{IJ} F^{+\mu\nu J} \tag{2.20}$$

And under a global symplectic transformation $\mathcal{S} \in \mathrm{Sp}(2n_v, \mathbb{R})$:

$$\begin{pmatrix} F'^+ \\ G'^+ \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} F^+ \\ G^+ \end{pmatrix} \implies \begin{aligned} F'^+ &= (A + B\mathcal{N})F^+ \\ G'^+ &= (C + D\mathcal{N})F^+ \end{aligned} \tag{2.21}$$

And, for consistency, we should get the same form $G'^+ = \mathcal{N}' F'^+$ after this transformation. Therefore one gets:

$$(C + D\mathcal{N})F^+ = \mathcal{N}'(A + B\mathcal{N})F^+ \tag{2.22}$$

This implies that, for coherence, the period matrix has to transform by a **Cayley (fractional-linear) map**,

$$\mathcal{N} \longmapsto (C + D\mathcal{N})(A + B\mathcal{N})^{-1}, \qquad \mathcal{S} = \begin{pmatrix} A & B \\ C & D \end{pmatrix}, \tag{2.23}$$

which, in general, produces a new kinetic matrix that cannot be written as $\mathcal{N}_{IJ}(\phi')$ for any field configuration ϕ' on the scalar manifold. In that case the Lagrangian genuinely changes form, and this is the microscopic origin of the inequivalence between symplectic frames. Only for those $\mathcal{S}$ that act on $\mathcal{N}$ in a way that can be reabsorbed into a diffeomorphism $\phi \to \phi'$ of the scalar manifold, the

Lagrangian remains invariant: these are precisely the isometries of $\mathcal{M}_{\text{scalar}}$, and they form the group G_U.

Only the electric vectors A^I_μ appear in the Lagrangian as independent degrees of freedom, while the magnetic duals $\tilde{A}_{\mu M}$ have no kinetic term, since they arise only through the dual field strengths G_I in the equations of motion. This asymmetry means that different choices of which n_v vectors to call "electric" lead to different Lagrangians. Each such choice is called a **symplectic frame**. To evaluate what is the role of these frames from a physical point of view, one must look separately at the ungauged and gauged theory. For the ungauged theory, choosing different symplectic frames leads to different, distinct Lagrangians related by $\mathrm{Sp}(2n_v, \mathbb{R})$ rotations outside G_U. At the level of the equations of motions and observables, these descriptions remain physically equivalent: they are simply the same system with different variables. When the theory is gauged via a subgroup $G_0 \subset G_U$, this equivalence no longer stands. Choosing a certain frame fixes the embedding of the gauge group generators inside the symplectic representation, with the consequence that different frames will yield genuinely different theories. Inequivalent frames for the full gauged deformation are classified by the double quotient

$$\mathrm{GL}(n_v, \mathbb{R}) \setminus \mathrm{Sp}(2n_v, \mathbb{R}) \,/\, G_U. \tag{2.24}$$

The $\mathrm{GL}(n_v, \mathbb{R})$ factor accounts for local field redefinitions of the n_v electric vectors, while G_U accounts for the freedom to redefine the scalar fields by an isometry. The choice of frame matters for the gauging procedure: it determines which generators of G_U can be realised as purely electric gauge symmetries. A different frame may allow gauging a group that would otherwise require magnetic vectors, or vice versa. In the presence of a duality group, the n_v electric vector fields A^M_μ and their magnetic partners $\tilde{A}_{\mu M}$ together form a $2n_v$-dimensional symplectic representation $\mathcal{R}_v$ of G_U. Since $\mathcal{R}_v$ is generally not the adjoint, the gauge fields now carry a symplectic index M rather than an adjoint index α, and one cannot use the standard covariant derivative (2.5) directly: one must specify which linear combination of generators $t_\alpha \in \mathfrak{g}_U$ each vector A^M_μ couples to, and different choices lead to physically inequivalent theories with different gauge groups $G_0 \subset G_U$, different matter couplings, and different vacuum structures [10].

The problem is therefore to identify a general, consistent gauging of $G_0 \subset G_U$ using vectors in $\mathcal{R}_v$, while maintaining full G_U-covariance of the description. The solution is the *embedding tensor*, introduced in the next section. Its peculiarity is that it preserves formal covariance under the full duality group G_U. By writing all defining elements of the gauged theory (couplings, covariant derivatives, variations) in terms of the G_U covariant embedding tensor, the Lagrangian preserves full covariance as well. Explicit symmetry breaking is then postponed to the moment in which one specific gauge group is chosen. This makes the problem of gauging classification significantly easier to handle, and in the following sections all advantages of the embedding tensor formalism will be concretely analyzed.

2.2 The embedding tensor

The central concept of this formalism is to encode the choice of gauge group in a single G_U-covariant tensor

$$\Theta_M^{\alpha} \in n_v \times \mathrm{adj}(G_U) \tag{2.25}$$

called embedding tensor. It can be viewed as a linear map from the vector representation $\mathcal{R}_v$ of G_U into the Lie algebra $\mathfrak{g}_U$:

$$\Theta : \mathcal{R}_v \to \mathfrak{g}_U , \quad e_M \mapsto \Theta_M{}^{\alpha}\, t_\alpha \,. \tag{2.26}$$

Where e_M is the M-th basis vector of the representation $\mathcal{R}_v$. This map selects which generator directions have nonzero coupling to the vector fields. The gauge generators that will enter the new covariant derivative are:

$$X_M = \Theta_M^{\alpha} t_\alpha \tag{2.27}$$

where M is the index of the fundamental representation, while α refers to the adjoint. The embedding tensor has a clear physical meaning: each row specifies which directions in the Lie algebra are coupled to the M-th vector field, and the nonzero entries in it mean that a certain generator t_α participates in the gauging, while null ones mean that a certain generator direction is not gauged. With these elements, it is possible to build the covariant derivative:

$$D_\mu \equiv \partial_\mu - g A_\mu^M X_M = \partial_\mu - g A_\mu^\Lambda\, \Theta_\Lambda{}^{\alpha}\, t_\alpha - g A_{\mu\Lambda}\, \Theta^{\Lambda\alpha}\, t_\alpha \,. \tag{2.28}$$

where g is the coupling constant and $A_\mu^M = (A_\mu^\Lambda, A_{\mu\Lambda})$ are all $2n_v$ (electric and magnetic) vector fields collectively. It is evident that the embedding tensor is the crucial object connecting the vector fields to the gauge symmetry: it selects a subalgebra $\mathfrak{g}_0 \subset \mathfrak{g}_U$ such that

$$\mathrm{rank}(\Theta_M{}^{\alpha}) = \dim(G_0) \,. \tag{2.29}$$

The rank of Θ counts the number of linearly independent gauge generators, but since they span precisely the gauge algebra $\mathfrak{g}_0$, the rank equals $\dim G_0$ by definition. The adjoint decomposes as $\mathfrak{g}_U = \mathfrak{g}_0 \oplus \mathfrak{g}_U/\mathfrak{g}_0$, where $\mathfrak{g}_0 = \mathrm{Im}(\Theta)$ contains the gauged generators and $\ker(\Theta^T)$ contains the ungauged (residual global) symmetries. In this sense, Θ acts as a *projector* that splits the global symmetry into a gauged and an ungauged part. Note also that $\dim(G_0)$ is bounded above by the number of available vector fields. In the electric frame only n_v vectors enter the Lagrangian, so

$$\dim(G_0) = \mathrm{rank}(\Theta) \leq n_v \,. \tag{2.30}$$

Intuitively, since each independent gauge direction requires a dedicated gauge boson to carry it, one cannot gauge more generators than there are vectors to act as connections. The rank of Θ as a matrix, therefore, corresponds precisely to the dimension of the gauge group. In the context of $\mathcal{N} = 8$ maximal supergravity, which will be presented in the next chapter, the embedding tensor $\Theta_M{}^{\alpha}$ is a $\mathbf{56} \times \mathbf{133}$ matrix (56 electric vectors, 133 generators of the exceptional group $\mathrm{E}_{7(7)}$). Different gaugings correspond to different ranks and rank structures of Θ, for example the SO(8) gauging corresponds to $\mathrm{rank}(\Theta) = 28$. Once $\Theta_M{}^{\alpha}$ is specified, the entire gauged Lagrangian is determined: the covariant derivatives, the non-abelian field strengths, the fermion shift matrices, and the scalar potential are all fixed functions of Θ. The embedding tensor encodes therefore the *complete data* of a gauging.

2.3 The scalar manifold and coset spaces

As mentioned above, the U-duality group G_U acts as the isometry group of the scalar manifold $\mathcal{M}_{scalar}$. This manifold can be characterized by generic, non-homogeneous geometries, but in many supergravity theories, especially highly constrained ones such as the $\mathcal{N} = 8$ maximal theory, it reduces to a homogeneous symmetric **coset space**. Therefore, in the following discussion, we will refer specifically to this setup.

To see more precisely the action of G_U on the scalar manifold, the concept of coset spaces needs to be properly introduced first. Given a Lie group G and a closed subgroup $H \subset G$, the **coset space** G/H is the set of equivalence classes

$$G/H = \{gH \mid g \in G\}, \qquad g \sim g' \quad \text{if} \quad g = g'h \quad \text{for} \quad g, g' \in G, \; h \in H. \tag{2.31}$$

of dimension $\dim G - \dim H$. If G has a transitive action on $\mathcal{M}_{\text{scal}}$ so that $g \cdot (g'H) = (gg')H$, the manifold is said to be *homogeneous*. This means that every point can be reached from any other point by a G transformation. H is called isotropy group, the subgroup of G that leaves a generic point on the manifold invariant. Homogeneity of scalar manifolds is very convenient since the coset structure can be used to obtain an algebraically tractable description of the symmetries involved in the gauging procedure. This framework provides the algebraic foundation to map the geometry of the scalar manifold into concrete matrix operators, allowing the construction of the Killing vectors $\xi^i_A(\phi)$ introduced in (2.18), and subsequently of the gauge couplings and the scalar potential.

At the Lie algebra level the structure can be decomposed into:

$$\mathfrak{g} = \mathfrak{h} \oplus \mathfrak{m}, \qquad \text{and} \qquad [\mathfrak{h}, \mathfrak{h}] \subset \mathfrak{h} \tag{2.32}$$

where $\mathfrak{h} = \text{Lie}(H)$ and $\mathfrak{m}$ is its complement in $\mathfrak{g}$. Coset manifolds can then be classified according to the structure of their algebra. For this purpose, a very useful concept is that of the Cartan–Killing metric

$$\eta_{AB} = f_{AD}{}^{C} f_{BC}{}^{D}. \tag{2.33}$$

where $f_{AD}{}^{C}$ are the structure constants of the full group G. The Cartan–Killing metric is non-degenerate for semi-simple groups, so in that case one can always map the generators to a basis where the metric is diagonal on $\mathfrak{g}$. In this basis, the coset manifold G/H is said to be **reductive**, and the additional condition $[\mathfrak{h}, \mathfrak{m}] \subset \mathfrak{m}$ holds: it states that $\mathfrak{m}$ transforms as a representation of H, which makes the tangent space to G/H well-defined. For a **symmetric coset manifold** the stronger condition

$$[\mathfrak{m}, \mathfrak{m}] \subset \mathfrak{h} \tag{2.34}$$

also holds. In supergravity, the geometry of the scalar manifold is severely constrained by supersymmetry: for theories with $\mathcal{N} > 2$ supercharges in $d = 4$, it is always a **symmetric space** [11]. Therefore, the scalar manifold is the coset manifold

$$\mathcal{M}_{\text{scalar}} = G_U/H, \tag{2.35}$$

where H is the maximal compact subgroup of G_U. The real dimension of this space equals the number of scalar fields in the theory.

A scalar field configuration is described by a local **coset representative** $\mathcal{L}(x) \in G_U$, a G_U-valued function of spacetime. Its transformation rule under rigid G_U and local H is

$$\mathcal{L}(x) \longmapsto g\,\mathcal{L}(x)\,h^{-1}(x), \qquad g \in G_U, \quad h(x) \in H. \tag{2.36}$$

The local H redundancy removes $\dim H$ unphysical degrees of freedom, leaving $\dim G_U - \dim H$ independent physical scalars. From $\mathcal{L}$ one constructs the **scalar matrix** in the symplectic representation,

$$\mathcal{M}_{MN}(\phi) = \left(\mathcal{L}(\phi)^T\,\mathcal{L}(\phi)\right)_{MN}, \tag{2.37}$$

which is symmetric, positive-definite, and H-invariant since the local h drops out as a consequence of H being the maximal compact subgroup. Under rigid G_U it transforms as $\mathcal{M} \to (g^{-1})^T \mathcal{M}\,g^{-1}$. It encodes the scalar kinetic term and, crucially, enters the scalar potential once the theory is gauged.

For $\mathcal{N} = 8$ supergravity, the relevant coset is $\mathrm{E}_{7(7)}/\mathrm{SU}(8)$, of real dimension $133 - 63 = 70$, corresponding to the 70 physical scalar fields of the theory. The embedding of the symplectic representation **56** of $\mathrm{E}_{7(7)}$ into $\mathrm{Sp}(56, \mathbb{R})$ will be discussed in Chapter 3.

2.4 Consistency and constraints

Not every embedding tensor is allowed: to build a consistent theory, it is necessary to impose constraints on Θ_M^α. This is another advantage of choosing the embedding tensor formalism: all consistency conditions for the choice of the gauge group become G_U-covariant constraints on Θ_M^α. In particular, three types of constraints must be satisfied, arising respectively from gauge invariance, locality, and supersymmetry.

1. Closure constraint

First, the gauge group must form an actual Lie algebra: the generators X_M must close under commutation, so they have to satisfy the relation:

$$[X_M, X_N] = f_{MN}^P X_P \tag{2.38}$$

where f_{MN}^P are the structure constants to be determined. Using (2.27) and the commutator for the generators t_α (2.2), the closure condition can be brought onto the embedding tensor, so that

$$\Theta_M{}^\alpha \Theta_N{}^\beta f_{\alpha\beta}{}^\gamma = f_{MN}{}^P \Theta_P{}^\gamma \tag{2.39}$$

This condition also implies the Jacobi identity for the structure constants of the gauge algebra, in the subspace projected by the embedding tensor:

$$f_{[MN}^Q f_{PQ]}^R \Theta_R{}^\alpha = 0 \tag{2.40}$$

Furthermore, the closure condition can be rewritten as

$$[X_M, X_N] = -X_{MN}{}^P X_P, \tag{2.41}$$

where $X_{MN}{}^P \equiv \Theta_M{}^\alpha (t_\alpha)_N{}^P$. This is the closure constraint written in terms of the $X_{MN}{}^P$ coefficients, that act as **generalised structure constants**: in fact, one could assume that they correspond to f^P_{MN} precisely, up to a sign; however, the matter is more complex than that. Consider the decomposition of X^P_{MN} into symmetric and antisymmetric parts:

$$X_{MN}{}^P = X_{[MN]}{}^P + X_{(MN)}{}^P \tag{2.42}$$

Only the antisymmetric part $X_{[MN]}{}^P$ plays the role of structure constants in the standard sense. The symmetric part $X_{(MN)}{}^P$ has no analogue in an ordinary Lie algebra and generates Chern–Simons-like couplings. $X_{(MN)}{}^P \equiv Z^P{}_{MN}$ can be isolated by defining a tensor orthogonal to $\Theta_M{}^\alpha$:

$$Z^{M\alpha}\Theta_M{}^\beta = 0 \qquad \Longleftrightarrow \qquad \begin{cases} Z^{\Lambda\alpha} = \frac{1}{2}\,\Theta^{\Lambda\alpha}\,, \\ Z_\Lambda{}^\alpha = -\frac{1}{2}\,\Theta_\Lambda{}^\alpha\,, \end{cases} \tag{2.43}$$

so that

$$X_{(MN)}{}^P = Z^{P\alpha}[t_\alpha]_M{}^Q\,\Omega_{NQ}\,. \tag{2.44}$$

It is now possible to introduce the non-abelian field strengths $\mathcal{F}_{\mu\nu}{}^M$. Using the usual expression $[D_\mu, D_\nu] = -g\mathcal{F}_{\mu\nu}{}^M X_M$, one could assume the form

$$\mathcal{F}_{\mu\nu}{}^M = \partial_\mu A^M_\nu - \partial_\nu A^M_\mu + g\,X_{[NP]}{}^M A^N_\mu A^P_\nu\,. \tag{2.45}$$

However, this is only true for the case of the electric frame, presented in the next section. In general, the variation of the field strengths transforms as

$$\delta_\xi \mathcal{F}_{\mu\nu}{}^M = -g\xi^P X_{PN}{}^M \mathcal{F}_{\mu\nu}{}^N + 2g\,Z^M{}_{PQ}\Big(\xi^P \mathcal{F}_{\mu\nu}{}^Q + A^P_{[\mu}\,\delta A^Q_{\nu]}\Big)\,, \tag{2.46}$$

with gauge parameter ξ^M, and it clearly has a non-covariant term depending on the presence of $Z^M{}_{PQ}$. The solution is the introduction of **2-form potentials** $B_{\mu\nu\,\alpha}$ in the adjoint of G_U, whose gauge transformation is tuned to cancel the Z-dependent term in (2.46). The improved field strength $\mathcal{H}_{\mu\nu}{}^M = \mathcal{F}_{\mu\nu}{}^M + Z^M_{PQ}B^{PQ}_{\mu\nu}$ then transforms covariantly. This procedure, known as **tensor hierarchy**, iterates: the B-field variation generates new non-covariant terms requiring 3-forms, and so on, but in $d = 4$ the hierarchy truncates at 4-forms.

Finally, an equivalent way to state the closure constraint is to require the embedding tensor to be *gauge-invariant*:

$$\Theta_M{}^\beta\,(t_\beta)_N{}^P\Theta_P{}^\alpha + \Theta_M{}^\beta\,f_{\beta\gamma}{}^\alpha\Theta_N{}^\gamma = 0, \tag{2.47}$$

since it must be a singlet under the action of the gauge group in order for the gauging to be self-consistent.

2. Locality constraint

The $2n_v$ vectors A^M_μ include both electric and magnetic potentials, but only mutually *local* vectors can simultaneously enter the gauge group. This is the locality condition: physically, it forbids the gauging from simultaneously using both an electric vector and its magnetic dual partner. The locality constraint reads

$$\Theta_M{}^\alpha\,\Theta_N{}^\beta\,\Omega^{MN} = 0 \qquad \Longleftrightarrow \qquad \Theta_I{}^{[\alpha}\Theta^{I\,\beta]} = 0, \tag{2.48}$$

where Ω^{MN} is the symplectic metric. This condition guarantees that there always exists a symplectic frame in which the gauging is purely electric.

3. Linear constraint

The embedding tensor $\Theta_M{}^\alpha$ transforms in the tensor product of the fundamental representation of the symplectic group (index M) and the adjoint of G_U (index α). This product generically decomposes into several irreducible representations. Supersymmetry, however, restricts $\Theta_M{}^\alpha$ to lie in a subset of these representations, precisely those whose contributions to the $O(g)$ supersymmetry variations can be cancelled, introducing by appropriate fermion-shift deformations of the Lagrangian. All other representations must be projected out. This is the **linear constraint**, which takes the schematic form

$$\mathbb{P}_{\text{forbidden}}\,\Theta = 0, \tag{2.49}$$

where $\mathbb{P}_{\text{forbidden}}$ is the projector onto the representations not allowed by supersymmetry.

As will be shown in the next chapter, for $\mathcal{N} = 8$ supergravity the U-duality group is $G_U = \mathrm{E}_{7(7)}$, with $\dim \mathfrak{e}_{7(7)} = 133$. The vector fields transform in the **56** of $\mathrm{E}_{7(7)}$, so $\Theta_M{}^\alpha$ sits in

$$\mathbf{56} \otimes \mathbf{133} = \mathbf{56} \oplus \mathbf{912} \oplus \mathbf{6480}. \tag{2.50}$$

Supersymmetry requires that only the **912** representation is retained:

$$\Theta_M{}^\alpha \;\in\; \mathbf{912} \;\subset\; \mathbf{56} \otimes \mathbf{133}. \tag{2.51}$$

The projection onto the **912** is equivalently expressed as

$$X_{(MNP)} = 0, \qquad X_{MNP} \equiv X_{MN}{}^Q\Omega_{QP}, \tag{2.52}$$

The 912 constraint (2.52) imposes the following symmetry conditions on $X_{MN}{}^P$:

$$X_{MN}{}^M = X_{MN}{}^N = 0\,, \qquad X_{M[NP]} = 0\,, \qquad X_{(MNP)} = 0\,, \tag{2.53}$$

where $X_{MNP} \equiv X_{MN}{}^Q\Omega_{QP}$.

The three constraints (2.47), (2.48) and (2.49) together form a complete, supersymmetric and gauge-consistent formulation of the theory. Note that (2.47) and (2.48) are both quadratic in Θ and together are often called the **quadratic constraints**. One must also note, however, that these constraints aren't entirely mutually independent: for any algebraically closed gauging that satisfies the linear supersymmetry constraint, locality is satisfied as well.

The scalar potential

A key consequence of the gauging is the generation of a **scalar potential** $V(\phi)$. For supergravity, in the ungauged theory ($g = 0$) there is no potential. Switching on the embedding tensor lifts this degeneracy by generating a potential entirely determined by Θ and the scalar matrix $\mathcal{M}_{MN}(\phi)$.

When considering a scalar manifold structured as the coset space $\mathcal{M}_{scalar} = G_U/H$, presented in Section 2.3, the potential is built through the **T-tensor** defined using the coset representative [11]:

$$\begin{aligned}\mathbb{T}(\Theta,\phi)_{\underline{M}} &= \mathcal{L}(\phi)^{-1}{}_{\underline{M}}{}^{N}\,\mathcal{L}(\phi)^{-1}X_N\mathcal{L}(\phi)\\ &= \mathcal{L}(\phi)^{-1}{}_{\underline{M}}{}^{N}\,\Theta_N{}^{\beta}\,\mathcal{L}(\phi)_{\beta}{}^{\alpha}\,t_\alpha\\ &= \mathbb{T}(\Theta,\phi)_{\underline{M}}{}^{\alpha}\,t_\alpha\,,\end{aligned} \tag{2.54}$$

where

$$\mathbb{T}(\Theta,\phi)_{\underline{M}}{}^{\alpha} = \mathcal{L}(\phi)^{-1}{}_{\underline{M}}{}^{N}\,\Theta_N{}^{\beta}\,\mathcal{L}(\phi)_{\beta}{}^{\alpha} = \left(L^{-1}(\phi)\star\Theta\right)_{\underline{M}}{}^{\alpha}\,, \tag{2.55}$$

where $\star$ denotes the action of L^{-1} as an element of G_U on Θ in the corresponding representation. Underlined indices label flat Lie-algebra components. The T-tensor decomposes under the local symmetry H into H-irreducible components called **fermion-shift tensors**. After gauging the Lagrangian, these tensors are needed to preserve supersymmetry: the minimal coupling $\partial \to D$ breaks the Ward identities of the ungauged supersymmetric theory at order g, and the fermion-shift tensors provide the terms that are necessary to compensate this effect. The **scalar potential** is a quadratic function of these shifts,

$$V(\phi) = g^2\Big[a\,\mathcal{M}^{MN}(\phi)\,\mathcal{M}^{PQ}(\phi)\,X_{MP}{}^{\alpha}X_{NQ\,\alpha} - b\left(\mathcal{M}^{MN}(\phi)\,X_{MN}{}^{\alpha}\right)^2\Big], \tag{2.56}$$

where $a, b > 0$ are numerical coefficients fixed by supersymmetry, and index contractions on α use the Killing form of $\mathfrak{g}_U$, $g_{\alpha\beta} = \mathrm{Tr}\big(\mathrm{ad}(t_\alpha)\ \mathrm{ad}(t_\beta)\big) = f_{\alpha\gamma}{}^{\delta}\,f_{\beta\delta}{}^{\gamma}$. The two terms arise from the square of the T-tensor and its H-trace respectively. In general, it is always possible to put the scalar potential in the form

$$V = g^2\,V^{MN}{}_{\alpha\beta}\,\Theta_M{}^{\alpha}\,\Theta_N{}^{\beta}\,, \tag{2.57}$$

where the dependence on the embedding tensor is manifest. $V^{MN}{}_{\alpha\beta}$ is a scalar-dependent matrix encoding the geometry of the scalar manifold. A **vacuum** of the gauged theory is a critical point ϕ_0 of this potential,

$$\nabla_i V|_{\phi_0} = 0, \tag{2.58}$$

and the value $V(\phi_0)$ determines the cosmological constant of the background: $V(\phi_0) > 0$ gives a de Sitter vacuum, $V(\phi_0) < 0$ an anti-de Sitter vacuum, and $V(\phi_0) = 0$ a Minkowski vacuum. One of the questions that this thesis wants to tackle is classifying consistent gaugings and their vacua. This, combined with the rank analysis that we want to conduct, reduces to the problem of finding embedding tensors that satisfy the constraints of the previous section, and analysing the structure of possible critical points.

2.5 The electric frame

To conduct the analysis on the embedding tensor's rank, presented in Chapter 4, the frame employed is purely electric, meaning that all gauged charges involved are electric. As a consequence of the locality constraint, this choice is perfectly admissible. For any gauged supergravity with duality group $\mathrm{G}_U \subset \mathrm{Sp}(2n_v, \mathbb{R})$, through a symplectic transformation, one can always switch between frames while maintaining generality (they give equivalent equations of motion, but different frames may correspond to different Lagrangians). Choosing this frame brings some additional simplification to the objects involved in the theory. A vector V^M in the $\mathbf{2n_v}$-dimensional symplectic representation can then be decomposed in electric and magnetic components in the following way:

$$V^M \longrightarrow (V^\Lambda,\, V_\Lambda) = (V^A,\, V^a,\, V_A,\, V_a), \tag{2.59}$$

with $A = 1, \ldots, r$ and $a = n_v - r + 1, \ldots, n_v$. Here electric (gauge field) components are denoted by upper indices A, a and their magnetic duals by corresponding lower indices A, a. The components V^a then span the subspace defined by the condition $\Theta^\alpha_\Lambda V^\Lambda = 0$. The first obvious consequence is that only the Θ^α_Λ and the $X_{MN}{}^P$ with $M = A$ are nonzero. Imposing the linear and quadratic constraints, the block decomposition of $X_{AN}{}^P$ (with row and column indices denoted by B, b and C, c respectively) takes the form:

$$X_{AN}{}^P = \begin{pmatrix} -f_{AB}{}^C & h_{AB}{}^C & C_{ABC} & C_{ABc} \\ 0 & 0 & C_{ACb} & 0 \\ 0 & 0 & f_{AC}{}^B & 0 \\ 0 & 0 & -h_{AC}{}^B & 0 \end{pmatrix}, \tag{2.60}$$

with the conditions

$$h_{(AB)}{}^C = C_{[AB]C} = C_{[ABC]} = f_{[AB]}{}^C = f_{AB}{}^B = 0. \tag{2.61}$$

In particular, the condition $f_{AB}{}^B = 0$ implies that the gauge group is unimodular. The closure relations (2.41) imply the identities,

$$\begin{aligned} f_{AB}{}^D f_{CD}{}^E &= 0, \\ f_{AB}{}^D h_{CD}{}^E + 4\, f_{[C|A}{}^E C_{B|D]e} &= 0, \\ f_{AB}{}^E C_{ECD} - 4\, f_{(C[A}{}^E C_{B]D)E} + 4\, h_{(C[A}{}^a C_{B]D)a} &= 0\,, \\ f_{AB}{}^D C_{C[De]} &= 0. \end{aligned} \tag{2.62}$$

These conditions can be used to study the case of embedding tensors having different ranks, especially lower ones like rank-1 and rank-2, since they provide meaningful simplification to the problem at hand.

The block decomposition (2.60) and the identities (2.61)–(2.62) represent the realization of the abstract embedding tensor constraints into explicit algebraic conditions on the tensors $f_{AB}{}^C$, $h_{AB}{}^C$, and C_{ABC} that label the gauging in the electric frame. They are the data one must supply to specify a gauging concretely: the structure constants $f_{AB}{}^C$ of the semisimple part, the torsion-like contribution $h_{AB}{}^C$, and the symmetric tensor C_{ABC} encoding the Chern–Simons couplings. Different choices of these tensors, all constrained by (2.61) and (2.62), correspond to inequivalent gauge algebras $\mathfrak{g}_0 \subset \mathfrak{g}_U$. One more thing to note is that $Z^P{}_{MN}$ does not vanish even in the electric frame, but does not appear in the Lagrangian since $\Theta^{\Lambda\alpha} = 0$, providing an additional simplification.

With this general framework in place, the next chapter specialises to $\mathcal{N} = 8$ maximal supergravity in four dimensions, where $\mathrm{G}_U = \mathrm{E}_{7(7)}$, $\mathrm{H} = \mathrm{SU}(8)$, the scalar manifold is $\mathrm{E}_{7(7)}/\mathrm{SU}(8)$, and the embedding tensor lives in the $\mathbf{912}$ of $\mathrm{E}_{7(7)}$. The power of the embedding tensor formalism is precisely that every result derived in this chapter carries over unchanged to that setting: one only needs to specify the group-theoretic data of $E_{7(7)}$ to obtain a fully determined, supersymmetric and consistent gauged theory.

Chapter 3

$\mathcal{N} = 8$ Supergravity in four dimensions

Among supergravity theories, the maximal $\mathcal{N} = 8$ theory is of unique significance: it is the four-dimensional theory with the maximum number of supercharges compatible with a massless graviton and no fields of spin greater than two. This characteristic was first established in 1978 by Cremmer and Julia first for the ungauged theory in [2] and then in the gauged SO(8) case [12], and later studied by de Wit and Nicolai in [13], with, finally, a full lagrangian formulation in terms of the embedding tensor provided by de Wit, Samtleben and Trigiante in [14] [4]. This means that the field content, the Lagrangian, and the symmetry structure are all completely fixed, so that no free choices have to be made. This unfortunately implies complexity, in the sense that the theory contains 256 on-shell degrees of freedom organised into a single supermultiplet, and its symmetry group is the exceptional Lie group $\mathrm{E}_{7(7)}$, which acts non-linearly on the 70 scalar fields parametrising the coset $\mathrm{E}_{7(7)}/\mathrm{SU}(8)$. In this chapter, the $\mathcal{N} = 8$ theory is introduced as the concrete setting for the application of the general embedding tensor formalism presented in Chapter 2. Starting from a review of the field content and the symmetry structure of the ungauged theory, I will present the coset $\mathrm{E}_{7(7)}/\mathrm{SU}(8)$ and the construction of the scalar matrix $\mathcal{M}_{MN}$ from the coset representative. $\Theta_M{}^\alpha$ and its constraints will be then adapted to the maximal case. Finally, I will present the gauged Lagrangian, the fermion-shift tensors and the scalar potential, and how its critical points emerge. The references used throughout this chapter are mainly [12], [13], [9].

3.1 The $\mathcal{N} = 8$ multiplet

Maximal supergravity in 4 dimensions preserves invariance under local supersymmetry transformations generated by N=8 spinorial charges. Knowing this, one can obtain the content of the theory. The resulting multiplet is presented in Table 3.1. This multiplet is CPT-self-conjugate [5], therefore the scalars result neutral with respect to the $U(1)$ subgroup of the $U(8)$ R-symmetry group. They transform non-trivially only with respect to the SU(8) subgroup. This group is then the holonomy group of the scalar manifold, that in this case is given by

$$\mathcal{M}^{\mathcal{N}=8}_{\text{scalar}} = \mathrm{E}_{7(7)}/\mathrm{SU}(8). \tag{3.1}$$

Therefore, as mentioned in the previous chapter, the duality group of maximal supergravity can be now identified as $\mathrm{G}_{\mathrm{U}} = \mathrm{E}_{7(7)}$ and its isotropy (and holonomy) group is $\mathrm{H} = \mathrm{SU}(8)$. Actually, there is an additional discrete symmetry $\mathbb{Z}_2$, so that the full H should be $\mathrm{SU}(8)/\mathbb{Z}_2$. However, for simplicity this detail will be omitted in writing since it is not relevant to this discussion.

Table 3.1: Field content of the $\mathcal{N} = 8$ supergravity multiplet. The 256 on-shell degrees of freedom are organized into five types of fields, all undergoing transformations under the R-symmetry group SU(8), whose properties are denoted by the properly antisymmetrized indices i, j, k, l. The Λ index for the vectors marks the duality representation under which they transform, maintaining notation from Chapter 2. The 70 real scalars span the coset space $\mathrm{E}_{7(7)}/\mathrm{SU}(8)$.

Spin	Physical state	Number	Representation	Field
2	Graviton	1	singlet	Vielbein e_μ^a
3/2	Gravitinos	8	**8** of SU(8)	Rarita-Schwinger ψ_μ^i
1	Vectors	28	**28** of SU(8)	Abelian gauge field A_μ^Λ
1/2	Fermions	56	**56** of SU(8)	Majorana spinors χ^{ijk}
0	Scalars	70	**70** of SU(8)	real scalar fields ϕ^{ijkl}

The content listed takes into account the independent degrees of freedom of the theory, leaving out the products of duality: the magnetic duals of the vector potentials and the 2-forms obtained from the dualization of the scalar fields. In the ungauged theory, the electric and magnetic potentials form a symplectic doublet that parametrize the same propagating degrees of freedom. Including both electric and magnetic potentials as dynamical fields in the Lagrangian would amount to double-counting the same degrees of freedom. In particular, only electric vector potentials will be present in the ungauged action. The global transformations (3.2) and (3.3) represent the action of the U-duality group $\mathrm{G_U} = \mathrm{E}_{7(7)}$ on the symplectic doublet of electric and magnetic vector potentials respectively, showing that for the magnetic case, the Lagrangian would change only for a pure derivative term.

$$\delta_\xi A_\mu^\Lambda = (t_\alpha)^\Lambda{}_N \xi^\alpha A_\mu^N \tag{3.2}$$

$$\delta_\xi A_{\mu\Lambda} = (t_\alpha)_\Lambda{}^N \xi^\alpha A_{\mu N} + \partial_\mu \Sigma_\Lambda(\xi) \tag{3.3}$$

where $(t_\alpha)^M{}_N$ are the 56×56 matrices representing the generators of $\mathrm{E}_{7(7)}$ in the symplectic representation. The term $\partial_\mu \Sigma_\Lambda(\xi)$ in (3.3) appears because a global $E_{7(7)}$ transformation rotates electric and magnetic fields, changing the Lagrangian overall by a total derivative. Therefore, to restore invariance, the magnetic potentials must undergo an induced gauge transformation, where $\Sigma_\Lambda(\xi)$ is a local function depending linearly on the global parameters ξ^α and the fields. When gauging is turned on (via the embedding tensor), the duality symmetry is partially broken: the electric charges are promoted to a local gauge symmetry, while magnetic charges are not. This asymmetry is why the Lagrangian contains only electric potentials; the magnetic potentials become non-dynamical or appear only through topological terms and 2-forms $B_{\mu\nu,\alpha}$. The non-covariant contributions from magnetic potentials in the formulas are cancelled by these 2-form gauge transformations via the tensor hierarchy, and appear only as total derivatives, thus decoupling from the action.

3.2 The scalar manifold $\mathrm{E}_{7(7)}/\mathrm{SU}(8)$

Together, the set of scalar fields describes a non linear σ-model given by the homogeneous manifold (3.1). To understand its structure, it is appropriate to present some of the properties of $\mathrm{E}_{7(7)}$: it is the split real form of the exceptional Lie group E_7, with $\dim \mathfrak{e}_{7(7)} = 133$, having rank 7. Its maximal compact subgroup is SU(8), with $\dim \mathfrak{su}(8) = 63$. Therefore, the coset has dimension $133 - 63 = 70$, that corresponds to the number of scalar fields in the multiplet. Having talked about cosets and relative Lie algebras in Chapter 2, applying this knowledge to this case, $\mathfrak{e}_{7(7)}$ can be decomposed as

$$\mathfrak{e}_{7(7)} = \mathfrak{su}(8) \oplus \mathbf{70}, \tag{3.4}$$

where $\mathfrak{su}(8)$ is the maximal compact subalgebra, while **70** denotes the 70 non-compact generators that transform in the antisymmetric self-dual representation of SU(8). The two subspaces satisfy the symmetric coset relations (2.34):

$$[\mathfrak{su}(8), \, \mathbf{70}] \subset \mathbf{70}, \qquad [\mathbf{70}, \, \mathbf{70}] \subset \mathfrak{su}(8). \tag{3.5}$$

The 70 scalar fields are described by the coset representative $\mathcal{L}(\phi) \in \mathrm{E}_{7(7)}$, whose general structure has been introduced in (2.36). It is defined up to a local SU(8) right transformation. Its Maurer–Cartan one-form, that is left-invariant, decomposes along two subspaces:

$$\mathcal{L}^{-1}\partial_\mu \mathcal{L} = Q_\mu + P_\mu, \tag{3.6}$$

Here, $Q_\mu \in \mathfrak{su}(8)$ is the composite SU(8) connection that carries the local H redundancy. $P_\mu \in \mathbf{70}$ is the coset vielbein. In the SU(8) index notation, $P_\mu = P_\mu^{ijkl}$ is totally antisymmetric, and satisfies the self-duality condition $P_\mu^{ijkl} = \frac{1}{4!}\varepsilon^{ijklmnpq}P_{\mu,mnpq}$. The number of resulting independent components is exactly 70, which is also the number of scalars fields. Considering $\mathcal{L}$ as a 56×56 symplectic matrix $L_M{}^N$, in the fundamental representation of $\mathrm{E}_{7(7)}$, one defines the symmetric positive-definite matrix

$$\mathcal{M}_{MN} = (L^T L)_{MN}, \tag{3.7}$$

The kinetic matrix $\mathcal{N}_{\Lambda\Sigma}$ that enters the kinetic term of the ungauged Lagrangian is obtained from the 28×28 blocks of $L_M{}^N$.

3.3 The ungauged Lagrangian

In order to have a complete picture, it is worth showing the form of the ungauged Lagrangian in maximal $\mathcal{N} = 8$ supergravity. This section does not aim to exhaust the topic, but only to provide the necessary basis to properly understand the gauged theory. As seen previously in this chapter, the Lagrangian will only enclose the actual fields dynamically active in the theory, leaving out unnecessary and redundant duals. In terms of the embedding tensor, the ungauged theory corresponds to setting $\Theta_M{}^\alpha = 0$. As seen in Chapter 2, this implies no need for covariant derivatives, and the field strengths are abelian, while the scalar potential (2.56) vanishes. Furthermore, the ungauged Lagrangian has a global symmetry within $\mathrm{E}_{7(7)}$, and a local SU(8) invariance. Following [9], let us look at the ungauged Lagrangian piece by piece, starting from the **bosonic sector**:

$$e^{-1}\mathcal{L}_{bos} = -\frac{1}{2}R - \frac{1}{2}G_{ij}(\partial_\mu\phi^i)(\partial^\mu\phi^j) - \frac{1}{4}\mathcal{M}_{MN}(\phi)F^M_{\mu\nu}F^{\mu\nu N} - \cdots \tag{3.8}$$

The first term is the Einstein-Hilbert term, and $e = \sqrt{-g}$. Since this is an abelian theory still, the vector fields appear in the Lagrangian through the field strengths:

$$F^M_{\mu\nu} = \partial_\mu A^M_\nu - \partial_\nu A^M_\mu \tag{3.9}$$

The terms in the dots correspond to the fermionic contributions and their higher-order interactions. G_{ij}, the scalar kinetic matrix, and $\mathcal{M}_{MN}(\phi)$, the analogue for vectors defined in (3.7), are the only objects left to determine, and their form is heavily constrained by supersymmetry. Using the tensors introduced in (3.6), the scalar term in the Lagrangian can be rewritten as

$$\mathcal{L}_{scalar} = -\frac{1}{2}Tr(P_\mu P^\mu). \tag{3.10}$$

However, a convenient choice would be expressing the scalar fields using the symmetric scalar matrix $\mathcal{M}_{MN}$ in order to keep the local SU(8) gauge freedom.

$$\mathcal{L}_{\text{scalar}} = \frac{1}{8}\,\text{Tr}\Big(\partial_\mu \mathcal{M}\,\partial^\mu \mathcal{M}^{-1}\Big) = \frac{1}{8}\,(\mathcal{M}^{-1})^{MN}\,\partial_\mu \mathcal{M}_{NP}\,(\mathcal{M}^{-1})^{PQ}\,\partial^\mu \mathcal{M}_{QM}. \tag{3.11}$$

As for the kinetic (bosonic) vector part of the Lagrangian, $\mathcal{L}_{vec} = -\frac{1}{4}\mathcal{M}_{MN}(\phi)F^M_{\mu\nu}F^{\mu\nu N}$, it is manifestly invariant under $\mathrm{E}_{7(7)}$. To analyse it along with the fermionic Lagrangian, I will use the notation developed in [4], that employs the self-dual and anti self-dual tensors $F^\Lambda_{\mu\nu} = F^{+\Lambda}_{\mu\nu} + F^{-\Lambda}_{\mu\nu}$ so that $\tilde{F}^{\pm,\Lambda}_{\mu\nu} = \pm F^{\pm\Lambda}_{\mu\nu}$.

$$\begin{aligned}\mathcal{L}_{\text{vector}} = -\frac{1}{4}(&iN_{\Lambda\Sigma}F^{+}_{\mu\nu\Lambda}F^{+\,\mu\nu\Sigma} - i\bar{N}_{\Lambda\Sigma}F^{-\Lambda}_{\mu\nu}F^{-\Lambda\mu\nu}) \\ &+ F^{+\Lambda}_{\mu\nu}O^{+\,\mu\nu}_{\Lambda} - F^{-\Lambda}_{\mu\nu}O^{-\,\mu\nu}_{\Lambda} \\ &+ i[(N-\bar{N})^{-1}]^{\Lambda\Sigma}[O^{+\mu\nu}_{\Lambda}O^{+}_{\mu\nu\Sigma} + O^{-\mu\nu}_{\Lambda}O^{-}_{\mu\nu\Sigma}]\end{aligned} \tag{3.12}$$

Here, $\mathcal{N}_{\Lambda\Sigma}$ is the field-dependent kinetic matrix, introduced in (2.16). Notice that the first term is just the vector bosonic sector rewritten in terms of the new field strengths. $O^{\pm\mu\nu}_{\Lambda}$ are bilinears of fermion fields, and terms quadratic in those terms are such that any additional terms in the Lagrangian, which no longer depends on the field strengths, transform covariantly under electric-magnetic duality. In terms of the gravitini ψ^i_μ and spin-$\frac{1}{2}$ fields χ^{ijk}, the bilinears $O^{\pm\mu\nu}_{\Lambda}$ are proportional to:

$$\mathcal{O}^{+\,ij}_{\mu\nu} = \frac{1}{2}\sqrt{2}\,\bar{\psi}^i_\rho\gamma^{[\rho}\gamma_{\mu\nu}\gamma^{\sigma]}\psi^j_\sigma - \frac{1}{2}\bar{\psi}_{\rho k}\gamma_{\mu\nu}\gamma^\rho\chi^{ijk} - \frac{1}{144}\sqrt{2}\,\varepsilon^{ijklmnpq}\bar{\chi}_{klm}\gamma_{\mu\nu}\chi_{npq} \tag{3.13}$$

$\mathcal{O}^{+}_{\mu\nu\Lambda}$ is related to the $SU(8)$ fermion bilinears $\mathcal{O}^{+ij}_{\mu\nu}$ via the scalar 56-vielbein acting as a bridge between the local $SU(8)$ and the global $E_{7(7)}$ structures. This relation is explicitly given by:

$$\mathcal{O}^{+}_{\mu\nu\Lambda} = \mathcal{L}_\Lambda{}^{ij}\mathcal{O}^{+ij}_{\mu\nu} + \mathcal{L}_{\Lambda ij}\mathcal{O}^{-ij}_{\mu\nu} \tag{3.14}$$

where $\mathcal{L}_\Lambda{}^{ij}$ and its conjugate $\mathcal{L}_{\Lambda ij}$ are the blocks of the coset representative matrix:

$$\mathcal{L}_M{}^{\underline{N}} = \Big(\mathcal{L}_M{}^{ij}, \mathcal{L}_{M\,kl}\Big) = \begin{pmatrix}\mathcal{L}_\Lambda{}^{ij} & \mathcal{L}_{\Lambda\,kl} \\ \mathcal{L}^{\Sigma\,ij} & \mathcal{L}^{\Sigma}{}_{kl}\end{pmatrix}. \tag{3.15}$$

that parametrize the scalar manifold. Now, putting all pieces together and substituting expressions appropriately, the final total Lagrangian is:

$$\begin{aligned}\mathcal{L} = &-\frac{1}{2}e\,R - \frac{1}{2}e^{\mu\nu\rho\sigma}\left(\bar{\psi}_\mu^{\,i}\gamma_\nu D_\rho\psi_{\sigma i} - \bar{\psi}_{\mu i}\overline{D}_\rho\gamma_\nu\psi_\sigma^{\,i}\right)\\ &-\frac{1}{4}ie\left\{N_{\Lambda\Sigma}F_{\mu\nu}^{+\,\Lambda}F^{+\,\mu\nu\Sigma} - \bar{N}_{\Lambda\Sigma}F_{\mu\nu}^{-\,\Lambda}F^{-\,\mu\nu\Sigma}\right\}\\ &-\frac{1}{12}e\left(\bar{\chi}^{ijk}\gamma^\mu D_\mu\chi_{ijk} - \bar{\chi}_{ijk}\gamma^\mu\overline{D}_\mu\chi^{ijk}\right) - \frac{1}{12}e\,|\mathcal{P}_\mu^{ijkl}|^2\\ &-\frac{1}{2}\sqrt{2}\,e\left\{\bar{\chi}_{ijk}\gamma^\nu\gamma^\mu\psi_\nu^{\,l}\left(\mathcal{P}_\mu^{ijkl} + \bar{\mathcal{P}}_\mu^{ijkl}\right) + \text{h.c.}\right\}\\ &+eF_{\mu\nu}^{+\,\Lambda}\mathcal{O}_\Lambda^{+\,\mu\nu} + eF_{\mu\nu}^{-\,\Lambda}\mathcal{O}_\Lambda^{-\,\mu\nu} - eL_{ij}^{\Lambda}L^{\Sigma ij}\left[\mathcal{O}_{\mu\nu\Lambda}^{+}\mathcal{O}_\Sigma^{+\,\mu\nu} + \mathcal{O}_{\mu\nu\Lambda}^{-}\mathcal{O}_\Sigma^{-\,\mu\nu}\right]\\ &+\mathcal{L}_4,\end{aligned} \tag{3.16}$$

and $\mathcal{L}_4$ contains four-fermion field terms not explicitly reported but that can be found in [4]. The covariant derivative on spinors is

$$D_\mu\psi_\nu^i = \partial_\mu\psi_\nu^i - \frac{1}{4}\omega_\mu^{ab}\gamma_{ab}\psi_\nu^i - (Q_\mu)^i_j\psi_\nu^j. \tag{3.17}$$

Here, the role of Q_μ, introduced in (3.6), as the connection within SU(8) is clear. In maximal supergravity, each piece of the Lagrangian is uniquely fixed by its field content and symmetry structure. The complete Lagrangian is invariant under the global $\mathrm{E}_{7(7)}$, up to total derivative terms, acting on the scalar coset and rotating the symplectic vector $(A_\mu^\Lambda, A_{\mu\Lambda})$ as in (3.2) and (3.3), and under the local SU(8) R-symmetry acting on all fermion indices.

The supersymmetry transformation rules of the ungauged theory are completely fixed by $\mathcal{N} = 8$ and take the form [4]:

$$\begin{aligned}\delta\psi_\mu^i = &\;2\mathcal{D}_\mu\epsilon^i + \frac{1}{4}\sqrt{2}\,\hat{F}_{\rho\sigma}^{-\,ij}\gamma^{\rho\sigma}\gamma_\mu\epsilon_j + \frac{1}{4}\bar{\chi}^{ikl}\gamma^a\chi_{jkl}\,\gamma_a\gamma_\mu\epsilon^j\\ &+\frac{1}{2}\sqrt{2}\,\bar{\psi}_{\mu k}\gamma^a\chi^{ijk}\gamma_a\epsilon_j - \frac{1}{576}\,\varepsilon^{ijklmnpq}\bar{\chi}_{klm}\gamma^{ab}\chi_{npq}\,\gamma_\mu\gamma_{ab}\epsilon_j\,,\end{aligned} \tag{3.18a}$$

$$\delta\chi^{ijk} = -2\sqrt{2}\,\hat{P}_\mu^{\,ijkl}\gamma^\mu\epsilon_l + \frac{3}{2}\,\hat{F}_{\mu\nu}^{-\,[ij}\gamma^{\mu\nu}\epsilon^{k]} - \frac{1}{24}\sqrt{2}\,\varepsilon^{ijklmnpq}\bar{\chi}_{lmn}\chi_{pqr}\,\epsilon^r\,, \tag{3.18b}$$

$$\delta e_\mu^{\,a} = \bar{\epsilon}^i\gamma^a\psi_{\mu i} + \bar{\epsilon}_i\gamma^a\psi_\mu^{\,i}\,, \tag{3.18c}$$

$$\delta L_M^{\,ij} = 2\sqrt{2}\,L_{Mkl}\left(\bar{\epsilon}^{[i}\chi^{jkl]} + \frac{1}{24}\,\varepsilon^{ijklmnpq}\bar{\epsilon}_m\chi_{npq}\right)\,, \tag{3.18d}$$

$$\delta A_\mu^{\,M} = -i\Omega^{MN}L_N^{\,ij}\left(\bar{\epsilon}^k\gamma_\mu\chi_{ijk} + 2\sqrt{2}\,\bar{\epsilon}_i\psi_{\mu j}\right) + \text{h.c.} \tag{3.18e}$$

where $\hat{F}_{\mu\nu}^{\pm\,ij}$ and $\hat{\mathcal{P}}_\mu^{\,ijkl}$ are SU(8)-covariantized versions of the field strength and of the coset vielbein which include additional fermionic terms to ensure covariance under supersymmetry transformations.

$$\hat{\mathcal{P}}_\mu^{\,ijkl} = \mathcal{P}_\mu^{\,ijkl} - \sqrt{2}\left(\bar{\psi}_\mu^{\,[i}\chi^{jkl]} + \frac{1}{24}\,\varepsilon^{ijklmnpq}\bar{\psi}_{\mu m}\chi_{npq}\right)\,, \tag{3.19a}$$

$$\hat{F}_{\mu\nu\,ij}^{+} = F_{\mu\nu\,ij}^{+} + \frac{1}{4}\bar{\psi}_\rho^{\,k}\gamma^\rho\gamma_{\mu\nu}\chi_{ijk} - \frac{1}{8}\sqrt{2}\,\bar{\psi}_i\left\{\gamma_{\mu\nu},\,\gamma^{\rho\sigma}\right\}\psi_{\sigma j}. \tag{3.19b}$$

3.4 The embedding tensor in $\mathcal{N}=8$

Before moving on to the gauged Lagrangian, let us examine the characteristics of the embedding tensor, which is the central object of the gauging. The general formalism presented in Chapter 2 now specialises to $\mathrm{G_U} = \mathrm{E}_{7(7)}$. The embedding tensor is a $\mathbf{56} \times \mathbf{133}$ matrix $\Theta_M{}^\alpha$, where $M \in \mathbf{56}$ (fundamental representation of $\mathrm{E}_{7(7)}$) and $\alpha \in \mathbf{133}$ (adjoint representation). In the electric frame, it reduces to a $\mathbf{28} \times \mathbf{133}$ matrix $\Theta_\Lambda{}^\alpha$, since $\Theta^{\Lambda\alpha} = 0$. The supersymmetry constraint $X_{(MNP)} = 0$ induces a requirement on its representation, so that the embedding tensor has to belong to the **912** block in $\mathrm{E}_{7(7)}$,

$$\Theta_M{}^\alpha \;\in\; \mathbf{912} \;\subset\; \mathbf{56} \otimes \mathbf{133} = \mathbf{56} \oplus \mathbf{912} \oplus \mathbf{6480} \tag{3.20}$$

This condition also implies the equations [4]:

$$(t^\alpha)^{MN}\Theta_{N\alpha} = 0, \qquad \text{and} \qquad (t_\beta t^\alpha)_M{}^N \Theta_N{}^\beta = -\tfrac{1}{2}\,\Theta_M{}^\alpha \tag{3.21}$$

Combined with the quadratic constraint (2.47) and locality (2.48), these are the only conditions Θ must satisfy. Under the maximal compact subgroup SU(8), the **912** decomposes as

$$\mathbf{912} \;\longrightarrow\; \mathbf{420} \oplus \overline{\mathbf{420}} \oplus \mathbf{36} \oplus \overline{\mathbf{36}}, \tag{3.22}$$

so the embedding tensor carries exactly these four irreducible pieces.

In the electric frame introduced in Chapter 2, Section 2.5, the nonzero components of Θ are $\Theta_\Lambda{}^\alpha$ $(\Lambda = 1, \dots, 28)$, and the gauge generators reduce to

$$X_\Lambda = \Theta_\Lambda{}^\alpha\, t_\alpha \;\in\; \mathfrak{e}_{7(7)}. \tag{3.23}$$

The gauge group G_0 is then the subgroup of $\mathrm{E}_{7(7)}$ generated by $\{X_\Lambda\}$, with $\dim \mathrm{G}_0 = \mathrm{rank}(\Theta) \leq 28$. Different choices of Θ in the **912** satisfying the quadratic constraint correspond to different gauge groups embedded in $\mathrm{E}_{7(7)}$. The SO(8) gauging presented in [13] corresponds to $\mathrm{rank}(\Theta) = 28$, with all 28 electric vectors gauging the 28-dimensional algebra $\mathfrak{so}(8)$.

3.5 Fermion shifts and the gauged Lagrangian

In the ungauged theory, the supersymmetry algebra closes on-shell on the equations of motion of the ungauged Lagrangian (3.16). A crucial consistency requirement, known as **Ward identity**, is that the variation of the action under supersymmetry vanishes, and in the ungauged theory this holds automatically. After gauging, this is no longer true. The minimal substitution $\partial_\mu \to D_\mu$ introduces extra terms of order g in the variations, which must be cancelled by adding **fermion-shift deformations** A_1^{ij}, $A_2^{ijk}{}_l$ to the transformation rules:

$$\delta_\epsilon \psi^i_\mu \;\longrightarrow\; 2D_\mu \epsilon^i + \sqrt{2}g\, A_1^{ij}\gamma_\mu \epsilon_j + \cdots, \tag{3.24}$$

$$\delta_\epsilon \chi^{ijk} \;\longrightarrow\; -2g\, A_2^{ijkl}\epsilon_l + \cdots, \tag{3.25}$$

To properly define the shift tensors A_1^{ij} and A_{2i}^{jkl}, it is necessary to formally introduce in this context the T-tensor, discussed generally in Section 2.4. Following the formulation in [4], the T-tensor is an $SU(8)$-covariant, field-dependent object obtained by dressing the constant embedding tensor

Θ with the scalar fields via the $E_{7(7)}/SU(8)$ coset representative $\mathcal{L}(x)$. In the fundamental **56** representation, it is defined as:

$$T_{\underline{M}}{}^{\underline{\alpha}}[\Theta, \phi] t_{\underline{\alpha}} \equiv \mathcal{L}^{-1}{}_{\underline{M}}{}^{N} \Theta_N{}^{\alpha} \left(\mathcal{L}^{-1} t_\alpha \mathcal{L} \right), \tag{3.26}$$

where the underlined indices refer to the local $SU(8)$ basis, and in the **56** representation this becomes:

$$T_{\underline{MN}}{}^{\underline{P}} = \mathcal{L}^{-1}{}_{\underline{M}}{}^{M} \mathcal{L}^{-1}{}_{\underline{N}}{}^{N} \mathcal{L}_P{}^{\underline{P}} X_{MN}{}^{P} \tag{3.27}$$

Since the T-tensor is obtained from the embedding tensor by a finite $\mathrm{E}_{7(7)}$ transformation [9], it lives in the same $\mathrm{E}_{7(7)}$-representation, inheriting from Θ the linear constraint

$$\mathbb{P} T = 0, \tag{3.28}$$

The **A-tensors** (or fermion-shift tensors), introduced first in Chapter 2, are precisely the H-irreducible components of the T-tensor. The tensor A_1^{ij}, the gravitini shift, is symmetric in i, j and transforms in the **36**, while A^i_{2jkl} is antisymmetric and traceless in jkl, transforms in the **420**. These shifts correspond to the **fermion mass matrices**.

$$A_1^{ij}(\phi) \sim \mathrm{proj}_{(\mathbf{36})} \mathbb{T}(\Theta, \phi), \qquad A^i_{2jkl}(\phi) \sim \mathrm{proj}_{(\mathbf{420})} \mathbb{T}(\Theta, \phi), \tag{3.29}$$

$$A_1^{ij} = \frac{4}{21} T_k^{ikj}, \qquad A_{2i}^{jkl} = -\frac{4}{3} T_i^{[jkl]}, \tag{3.30}$$

By turning around this perspective, one can see the origin of the linear constraint (2.49): the supersymmetry-violating terms introduced by gauging are proportional to $\Theta_M{}^{\alpha}$ and must be cancelled by the variation of fermionic mass terms, and this is possible if and only if Θ lies in a representation compatible with the available fermionic bilinears $(\bar{\psi}\psi)$, $(\bar{\psi}\chi)$, $(\bar{\chi}\chi)$, which in $\mathcal{N} = 8$ fall into the $\mathbf{36} + \mathbf{420}$ + h.c. of SU(8), which is the content of the **912** representation. If the embedding tensor were to lie outside this representation, it would generate unavoidable terms, therefore it must be $\mathbb{P}\,\Theta = 0$. The introduction of the A-tensors is the only addition to restore supersymmetry in the gauged Lagrangian, apart from the fix of the covariant derivative via the substitution $\mathcal{F}_{\mu\nu} \to \mathcal{H}_{\mu\nu}$ discussed in the previous chapter as well.

Moving on to the gauged Lagrangian's expression, the introduction of a covariant derivative and the presence of a non-zero embedding tensor has now caused terms of order g and g^2 to appear. While the order g terms will generate mass-like terms for fermions, contained in $\mathcal{L}_{\text{shifts}}$, the newly generated potential $V(\phi)$ will be proportional to g^2:

$$\mathcal{L}_{\text{gauged}} = \mathcal{L}_{\text{ungauged}}\big|_{\partial \to D} + \mathcal{L}_{\text{C-S}} + \mathcal{L}_{\text{shifts}} - e\, V(\phi), \tag{3.31}$$

$\mathcal{L}_{\text{C-S}}$ contains the topological Chern–Simons couplings generated by the tensor C_{ABC} in (2.60). These tensors enter the fermion-shift terms as [4]:

$$\mathcal{L}_{\text{masslike}} = eg \left\{ \frac{1}{2}\sqrt{2}\, A_{1\,ij}\, \bar{\psi}^i_\mu \gamma^{\mu\nu} \psi^j_\nu + \frac{1}{6} A_{2i}^{jkl}\, \bar{\psi}^i_\mu \gamma^\mu \chi_{jkl} + A_3^{ijk,lmn}\, \bar{\chi}_{ijk} \chi_{lmn} \right\} + \text{h.c.}\,. \tag{3.32}$$

While the scalar potential follows the general form presented in (2.57) but now with the appropriate coefficients:

$$V(\phi) = g^2\left(\frac{1}{24}\, A_{2,i}{}^{jkl}\, A^i_{2jkl} - \frac{3}{4}\, A_1^{ij}\, A_{1ij}\right). \tag{3.33}$$

As for what identifies as a **vacuum** of the potential [11], one must remember that a Lorentz-preserving vacuum is a maximally symmetric solution in which all scalar fields take a constant value $\phi_0 \equiv \langle \phi^i(x) \rangle$, with other fields (vectors and fermions) vanishing on the solution. This v.e.v. corresponds to a critical point of the scalar potential,

$$\nabla_i V|_{\phi_0} = 0. \tag{3.34}$$

and can be categorized as Minkowski, De Sitter or Anti-De Sitter. The existence and nature of the different vacua depend on which gauge group $G_0 \subset E_{7(7)}$ is chosen. Different gaugings can admit different types of vacua, and a given gauging may admit multiple critical points corresponding to different phases of the theory. Once one has computed the expression for the first and second derivatives of the potential as a function of the embedding tensor, the search for vacua can start. However, to approach the problem in a convenient way, one can use the properties of coset spaces. Given that the embedding tensor transforms in the **912** of $\mathrm{E}_{7(7)}$, one can switch between different embeddings of G_0 through an $\mathrm{E}_{7(7)}$ transformation. Since $\mathrm{E}_{7(7)}$ is also the isometry group of the scalar manifold, the transformation will act not only on the gauge generators but also on the coset representatives, so the T-tensor will transform as:

$$T_{MN}{}^P = L^{-1}{}_M{}^{\bar{M}}\, L^{-1}{}_N{}^{\bar{N}}\, X_{\bar{M}\bar{N}}{}^{\bar{P}}\, L_{\bar{P}}{}^P, \xrightarrow{U} T'_{MN}{}^P = (U^{-1}L)^{-1}{}_M{}^{\bar{M}}\, (U^{-1}L)^{-1}{}_N{}^{\bar{N}}\, X_{\bar{M}\bar{N}}{}^{\bar{P}}\, (U^{-1}L)_{\bar{P}}{}^P, \tag{3.35}$$

$$X_{MN}{}^P = \Theta_M{}^\alpha\, (t_\alpha)_N{}^P \;\longrightarrow\; X'_{MN}{}^P = U_M{}^Q\, U_N{}^R\, X_{QR}{}^S\, U^{-1}{}_S{}^P, \tag{3.36}$$

where $U \in E_{7(7)}$. From a physical point of view, this means that a vacuum at a generic scalar v.e.v. ϕ_0 of a theory defined by Θ is equivalent, up to a global $\mathrm{E}_{7(7)}$ rotation, to a vacuum at the origin $\phi = 0$ of a theory with rotated embedding tensor $\Theta' = U \cdot \Theta$, without loss of generality. In this way, all quantities that characterise the gauged theory at the vacuum, such as fermion shifts, mass matrices, residual supersymmetry, can be computed at the origin. Basically, one can look for vacua corresponding to any value of the scalar fields by always staying in the origin, simply changing the embedding.

Therefore, scanning for vacua reduces to the algebraic problem of classifying embedding tensors in the **912** modulo $\mathrm{E}_{7(7)}$ equivalence. At the origin, the stationarity condition $D_x V = 0$ becomes an algebraic constraint on Θ, making it tractable for a systematic search.

3.5.1 Note on vacua and the electric frame

Choosing the electric frame restricts the embedding tensor components by the condition $\Theta^{\Lambda\alpha} = 0$. Consequently, only these 28 electric vectors mediate the gauging G_0. However, a generic choice of G_0 might possess a profile for the embedding tensor that does not respect this specific split, meaning that one gets $\Theta^{\Lambda\alpha} \neq 0$ in the initial reference frame. Instead of performing a duality rotation to find a purely electric frame, the embedding tensor formalism preserves full covariance. When non-vanishing magnetic components $\Theta^{\Lambda\alpha}$ are unavoidable, the naive Lagrangian is no longer gauge-invariant, and consistency requires introducing 2-form tensor fields $B_{\mu\nu\,\alpha}$ via the tensor hierarchy to compensate, as already noted in Section 3.3.

A crucial observation concerns the role of the symplectic group $\mathrm{Sp}(56, \mathbb{R}) \supset \mathrm{E}_{7(7)}$. For the ungauged theory all symplectic frames are physically equivalent: the equations of motion are $\mathrm{Sp}(56, \mathbb{R})$-covariant and any subset of the abelian vectors can be dualized freely, while $\mathrm{E}_{7(7)}$ acts as the on-shell duality group (a symmetry of the field equations, not of the action, whose off-shell part is only the electric subgroup of a given frame). Gauging breaks this structure: once a gauge group is switched on, the non-abelian vectors can no longer be dualized, and the choice of which 28 vectors are electric becomes physical. The only residual freedom relating frames is then $\mathrm{E}_{7(7)}$ together with field redefinitions. Symplectic rotations lying in $\mathrm{Sp}(56, \mathbb{R}) \setminus \mathrm{E}_{7(7)}$ cannot be reabsorbed and map a gauging to a genuinely inequivalent one. Hence two embedding tensors Θ and Θ' gauging the same abstract group G_0 but in different symplectic frames can give physically inequivalent supergravities with different scalar potentials, vacua, and mass spectra.

The canonical example is the **dyonic SO(8) gauging** [15], [16]: the standard de Wit–Nicolai gauging [13] is purely electric, with all 28 vectors gauging $\mathfrak{so}(8)$. Applying a symplectic rotation by an angle $\omega \in [0, \pi/8]$, which lies outside $\mathrm{E}_{7(7)}$, produces a one-parameter family of dyonic SO(8) gaugings with both electric and magnetic components non-zero. These are physically inequivalent: different values of ω yield different vacuum landscape and different mass spectra, even though the gauge group remains SO(8). This shows concretely that the classification of gaugings requires specifying not just G_0, but also its embedding in $\mathrm{Sp}(56, \mathbb{R})$. The choice of symplectic frame contains physical data, and is not merely a choice of variables.

Chapter 4

Scanning gaugings: the algorithm

The theoretical framework established in the previous chapters for maximal $\mathcal{N} = 8$ supergravity explained that a gauging is entirely specified by the embedding tensor $\Theta_M{}^\alpha$, lying in the **912** representation of $E_{7(7)}$. It has to satisfy the linear and quadratic constraints presented in Section 2.4 of Chapter 2. The problem of classifying gaugings (and, subsequently, vacua) becomes then the algebraic problem of finding all embedding tensors that satisfy those constraints, up to $E_{7(7)}$ equivalence. This chapter describes the systematic approaches used to try to solve this problem, looking at it specifically from the perspective of **rank stratification**. Instead of scanning the full variety of embedding tensors at once, one fixes the rank $r = \mathrm{rank}(\Theta)$ and searches for valid gaugings for each value $1, 2, \ldots, 28$. In particular, we aim to identify the lowest rank that admits any gauging at all. This analysis is conducted at first in the canonical electric frame, specifically in the electric SL(8) frame, where each choice of r vectors automatically satisfies the locality constraint. By imposing the supersymmetry constraint, a linear system in the components of Θ will be left to solve. The analysis of this system will lead to possible solutions that will then be filtered for closure checks using two approaches: a first attempt at classification involved the use of stochastic methods, but there were several limitations to the approach led to the formulation of an exact method of computation, which also showed limitations but filled some missing information. The gaugings found are then classified by calculating the appropriate invariants and, when feasible, the aid of the LieART [17] package. While for the **rank-1** and **rank-2** cases it is possible to reach certain conclusions analytically, for higher ranks one needs to utilize computational tools. Furthermore, the results obtained are of course only valid for this specific frame and, for quantities that are orbit-invariants, up to rotations inside $E_{7(7)}$, which stay in the same orbit. To prove that the result obtained for the lowest possible rank that admits any gauging is genuinely invariant, we provide strong computational evidence using an optimized frame-independent calculation of the minimal support admitted by the linear condition (2.49). The code written for this thesis was implemented mainly using Wolfram Mathematica and partially in Python.

4.1 The strategy

In this section, we will present the strategy chosen to conduct the rank stratification study from an analytical point of view, then we will move onto the practical aspects of its adaptation in the implementation of the code.

4.1.1 Ansatz and isotropy condition

The choice of classifying gaugings by the embedding tensor's rank is motivated by its link to the dimension of the gauge group G_0:

$$\mathrm{rank}(\Theta_M{}^\alpha) = \dim(\mathrm{G}_0)\,. \tag{4.1}$$

Now, in the case of $\mathcal{N} = 8$ supergravity, one has 56 vectors, however only 28 will enter the Lagrangian due to duality. Therefore, the biggest dimension G_0 can have is 28, and our rank search can stop at this number. Since the embedding tensor is a $\mathbf{56} \times \mathbf{133}$ object, we apply the following ansatz to rewrite its form:

$$\Theta = V^T C, \qquad V \in \mathbb{R}^{n\times 56}, \quad C \in \mathbb{R}^{133\times n}, \tag{4.2}$$

$$\Theta^\alpha_M = \sum_{a=1}^{n} (V_a)_M C^\alpha_a, \tag{4.3}$$

where the rows of V are the vectors V_a and C are the coefficient vectors C_a. Of course, $a = 1, ..., n$ where $n = \mathrm{rank}(\Theta)$, and $M = 1, ..., 56$ and $\alpha = 1, ..., 133$ are the usual fundamental and adjoint indices. The singular term with fixed a, $(V_a)_M C^\alpha_a$ is rank-1, so the sum of n terms has at most rank n, since:

$$\mathrm{rank}(\Theta) = \mathrm{rank}(V^T C) \leq \min(\mathrm{rank}(V), \mathrm{rank}(C)) \leq n. \tag{4.4}$$

The rank is exactly n when both sets $\{V_a\}$ and $\{C^\alpha_a\}$ are linearly independent. The ansatz of n independent vectors is, therefore, precisely the search for gauge groups of dimension n.

The frame of choice for this analysis is the **canonical electric frame**, keeping the formalism from Chapter 2, Section 2.5, and using unit vectors as a basis. This means that only the first 28 entries of V^M_a can be nonzero, and will take the form:

$$V_a = e_a \in \mathbb{R}^{56}, \qquad (e_a)_M = \delta_{Ma}, \qquad a = 1, \ldots, n \leq 28, \tag{4.5}$$

and the embedding tensor is

$$\Theta_M{}^\alpha = \sum_{a=1}^{n} \delta_{Ma}\, C_a{}^\alpha = C_M{}^\alpha, \qquad M = 1, \ldots, n, \tag{4.6}$$

$$\Theta_M{}^\alpha = \begin{pmatrix} C_1{}^\alpha \\ \vdots \\ C_n{}^\alpha \\ 0 \\ \vdots \\ 0 \end{pmatrix} \in \mathbb{R}^{56\times 133}, \qquad \text{rows } n+1, \ldots, 56 \text{ vanish.} \tag{4.7}$$

This choice is particularly well-suited for two reasons, the first being of course its clean and simple form, and the second being the fact that, this way, locality is always automatically ensured. Let us see what the locality constraint looks like now, with the substitution of the ansatz (4.3) into it:

$$\Omega^{MN}\Theta_M{}^\alpha\Theta_N{}^\beta = 0 \Longrightarrow \sum_{a,b=1}^{n} C_a{}^\alpha\, C_b{}^\beta\, \Omega^{MN} (V_a)_M (V_b)_N = 0 \qquad \forall\, \alpha, \beta. \tag{4.8}$$

Choosing the electric frame means that this condition is satisfied independently of the values of the coefficients C_a^α, so that locality is decoupled entirely from the supersymmetry constraint, and the problem will reduce to a linear system only. Imposing that the locality condition must hold for every valid solution C_a^α, not only for special combinations of them, makes the problem a decoupled linear one, instead of being coupled and nonlinear. Locality is therefore expressed as:

$$\Omega(V_a, V_b) \equiv \Omega^{MN}(V_a)_M(V_b)_N = 0 \qquad \forall\, a, b = 1, \dots, n, \tag{4.9}$$

which corresponds to saying that the n vectors V_a span an **isotropic subspace** of $(\mathbb{R}^{56}, \Omega)$. For a 56-symplectic space, this subspace can have at most dimension 28, which is consistent with the rank bound for Θ. The symplectic form Ω is a block matrix that has the top-left 28×28 block null, therefore it is evident that $\Omega^{ab} = 0$ for $a, b \leq 28$. Intuitively, the form ties each vector to its dual, but in the canonical electric frame, there are no magnetic duals active, therefore the result is always zero.

$$\Omega^{MN} = \begin{pmatrix} 0 & -\mathbf{1}_{28} \\ \mathbf{1}_{28} & 0 \end{pmatrix}. \tag{4.10}$$

Substituting the canonical unit vectors e_a:

$$\Omega(e_a, e_b) = 0, \qquad a, b = 1, ..., n. \tag{4.11}$$

Note that the locality constraint implies that there is always a choice of a symplectic frame so that the gauging can be made purely electric [5]. However, conclusions made for one electric frame can only be extended to frames that are in its same orbit, while there are other gaugings in different orbits that never touch the specific electric frame of choice. This is why, later in the chapter, a study that aims to extend some of the conclusions in this thesis to all orbits of $E_7(7)$ is built.

4.2 The supersymmetry constraint

The central condition that will be used to constrain the remaining coefficients C_a^α is the supersymmetry condition (2.52):

$$X_{(MNP)} = 0 \tag{4.12}$$

Starting from the generalized structure constants, substituting the rank-n ansatz and then lowering indices via Ω_{MN}, one gets:

$$X_{MN}{}^P \equiv \Theta_M{}^\alpha\,(t_\alpha)_N{}^P = \sum_{a=1}^{n} \delta_{Ma}\, C_a{}^\alpha\,(t_\alpha)_N{}^P, \qquad M = 1, \dots, n. \tag{4.13}$$

$$X_{MNP} \equiv X_{MN}{}^Q\,\Omega_{QP} = \sum_{a=1}^{n} \delta_{Ma}\, C_a{}^\alpha\,(t_\alpha)_N{}^Q \Omega_{QP}. \tag{4.14}$$

Symmetrizing the indices gives back the supersymmetry constraint as a linear system in $C_a{}^\alpha$. One can contract this with a set of test triples $(v, w, u) \in (\mathbb{R}^{56})^3$, which yields one linear equation in the $133n$ unknowns $(C_1{}^\alpha, \dots, C_n{}^\alpha)$:

$$v^M w^N u^P\, X_{(MNP)} = 0 \iff v^M w^N u^P X_{MN}{}^Q \Omega_{QP} + \mathrm{cyclic}(v, w, u) = 0, \tag{4.15}$$

$$\sum_{a=1}^{n}\sum_{\alpha=1}^{133}\Big[v^a\, w^N(\Omega\cdot t_\alpha)_{NP}\, u^P + w^a\, v^N(\Omega\cdot t_\alpha)_{NP}\, u^P + u^a\, v^N(\Omega\cdot t_\alpha)_{NP}\, w^P\Big] C_a{}^\alpha = 0, \tag{4.16}$$

where $v^a \equiv v^M \delta_{Ma}$ denotes the a-th component of v, and $(\Omega\cdot t_\alpha)_{NP} \equiv \Omega_{QP}(t_\alpha)_N{}^Q$ are 56×56 matrices. Assembling one row per test triple gives the linear system

$$\begin{aligned} A\,\mathbf{c} &= \mathbf{0}, \qquad A \in \mathbb{R}^{N_{\text{rows}}\times 133n},\\ \mathbf{c} &= (C_1{}^1,\ldots,C_1{}^{133},\ldots,\,C_n{}^1,\ldots,C_n{}^{133})^T \in \mathbb{R}^{133n}, \end{aligned} \tag{4.17}$$

where N_{rows} is the number of test triples used, so each one defines a row of the constraint matrix $\mathbf{A}$.

4.3 Analysing the null space

Once the linear system has been identified, finding solutions is basically a matter of identifying the null space of $\mathbf{A}$ and its dimension, since solutions will lie in $\ker(\mathbf{A})$. For the lowest rank-1 case, (4.17) can be managed manually, finding conditions one by one, for each row, that exhaust the problem, as will be shown in the next section. However, for higher ranks, it is necessary to employ computational methods, explicitly inserting the generators for $\mathrm{E}_{7(7)}$ and the metric Ω_{MN}, and finding the resulting null space. If its dimension is zero, of course the system will not admit any solutions; otherwise, the code will move onto the next steps of checking whether the null vectors satisfy the closure relation (2.41). In fact, an adapted closure condition can be constructed from the null vectors $G_a = C_a{}^\alpha t_\alpha$, starting from the canonical one:

$$[X_M, X_N] = -X_{MN}{}^P X_P, \tag{4.18}$$

where the ansatz for the embedding tensor will be plugged into X_M and $X_{MN}{}^P$. As already stated, in the canonical electric frame with rank-n gauge algebra, the embedding tensor is identified with the null-space coefficients:

$$\Theta_M{}^\alpha = C_M{}^\alpha, \qquad M = 1,\ldots,n, \tag{4.19}$$

where $C_M{}^\alpha$ are the coefficient vectors extracted from the null space. The gauge generators are constructed as:

$$X_M = \Theta_M{}^\alpha t_\alpha = C_M{}^\alpha t_\alpha, \qquad M = 1,\ldots,n. \tag{4.20}$$

The generalised structure constants (coefficients in the embedding tensor formalism) are:

$$X_{MN}{}^P = \Theta_M{}^\alpha (t_\alpha)_N{}^P = C_M{}^\alpha (t_\alpha)_N{}^P. \tag{4.21}$$

The standard closure relation becomes:

$$[C_M{}^\alpha t_\alpha, C_N{}^\beta t_\beta] = C_M{}^\alpha C_N{}^\beta [t_\alpha, t_\beta] = C_M{}^\alpha C_N{}^\beta f_{\alpha\beta}{}^\gamma t_\gamma. \tag{4.22}$$

In the computational framework, we work with the n-dimensional subalgebra of gauge generators (relabelling $M \to a$ for clarity):

$$G_a = C_a{}^\alpha t_\alpha, \qquad a = 1,\ldots,n, \tag{4.23}$$

so that the final closure condition to check for is:

$$[G_a, G_b] = C_a{}^\alpha C_b{}^\beta f_{\alpha\beta}{}^\gamma t_\gamma = f_{ab}{}^c\, C_c{}^\gamma\, t_\gamma = f_{ab}{}^c G_c, \tag{4.24}$$

where $f_{ab}{}^c$ are the structure constants of the rank-n gauge algebra, extracted via:

$$f_{ab}{}^c = \left(C^T\right)^+{}_c{}^\gamma\, (C_a{}^\alpha C_b{}^\beta f_{\alpha\beta}{}^\gamma), \tag{4.25}$$

where $\left(C^T\right)^+$ is the Moore–Penrose pseudoinverse [18], a generalization of the inverse matrix used often to optimize the resolution of linear systems. Given n generators $G_a = C_a{}^\alpha t_\alpha$ of a candidate gauge subalgebra, the commutator $[G_a, G_b] = C_a{}^\alpha C_b{}^\beta f_{\alpha\beta}{}^\gamma t_\gamma$ is a vector in the full 133-dimensional adjoint representation of $\mathrm{E}_{7(7)}$. If the n generators genuinely close into a Lie algebra, this commutator must be expressible as a linear combination of the generators themselves, so that $[G_a, G_b] = f_{ab}{}^c G_c$ for some structure constants $f_{ab}{}^c$. Comparing coefficients of t_γ on both sides yields the linear system

$$f_{ab}{}^c\, C_c{}^\gamma = C_a{}^\alpha C_b{}^\beta f_{\alpha\beta}{}^\gamma, \tag{4.26}$$

which has n unknowns ($f_{ab}{}^c$ for fixed a, b) but 133 equations, one for each generator direction in $\mathfrak{e}_{7(7)}$. Since this system is overdetermined, a regular matrix inverse does not exist. The Moore–Penrose pseudoinverse $(C^T)^+$ provides the unique minimum-norm least-squares solution: among all possible $f_{ab}{}^c$, it selects the one whose corresponding linear combination of generators best reproduces the commutator.

4.4 Analytical demonstrations for rank-1 and rank-2

In this section we will present the analytic resolution of the lowest ranks 1 and 2. It will provide meaningful results and serve as a practical example to understand the application of the strategy explained until now, and to have an idea of what the system to be solved looks like. It will also become clear that the complexity of this analysis grows rapidly with rank, and these are in fact the only two cases that are worth trying to solve manually, instead of relying fully on computational methods.

4.4.1 Rank-1

The singular-rank case is simple enough that it can be treated analytically, without the need to employ the code. I aim to demonstrate that the only solution compatible with the linear constraint $X_{(MNP)} = 0$ is

$$\Theta_M{}^\alpha = 0. \tag{4.27}$$

Hence, the rank-1 case gives only the ungauged, physically trivial theory.

Proof. Assume that the embedding tensor takes the rank-1 form:

$$\Theta_M{}^\alpha = V_M C^\alpha, \tag{4.28}$$

for some vector V_M in the **56** and some adjoint vector C^α, following the formalism presented in the previous sections. Define

$$\mathcal{C}_N{}^P := C^\alpha (t_\alpha)_N{}^P. \tag{4.29}$$

Then

$$X_{MN}{}^P = V_M \mathcal{C}_N{}^P. \tag{4.30}$$

Lowering the last index with the symplectic form,

$$\mathcal{C}_{NP} := \mathcal{C}_N{}^Q \Omega_{QP}, \qquad X_{MNP} = V_M \mathcal{C}_{NP}. \tag{4.31}$$

Since the generators t_α are symplectic, it follows that

$$(t_\alpha)_N{}^Q \Omega_{QP} = (t_\alpha)_P{}^Q \Omega_{QN}, \tag{4.32}$$

so each matrix $(t_\alpha)_{NP} := (t_\alpha)_N{}^Q \Omega_{QP}$ is symmetric. This symmetry transfers to $\mathcal{C}_{NP}$ so that $\mathcal{C}_{NP} = \mathcal{C}_{PN}$. The next step is to substitute the embedding tensor's expression into each independent constraint. First, putting the rank-1 ansatz into the locality constraint $\Theta_M{}^\alpha \Theta_N{}^\beta\, \Omega^{MN} = 0$, gives

$$V_M C^\alpha\, V_N C^\beta\, \Omega^{MN} = 0, \tag{4.33}$$

and reordering

$$C^\alpha C^\beta V_M V_N \Omega^{MN} = 0. \tag{4.34}$$

However, given that $V_M V_N$ is symmetric in (M, N), while Ω^{MN} is antisymmetric,

$$V_M V_N \Omega^{MN} = 0 \tag{4.35}$$

identically. Thus the locality constraint imposes no extra restriction in the rank-1 case, and it is automatically satisfied even before choosing a specific frame. This constraint provides no new useful information.

Moving on to the quadratic constraint, substituting $\Theta_M{}^\alpha$ into (2.47), gives:

$$V_M C^\beta (t_\beta)_N{}^P V_P C^\alpha + V_M V_N C^\beta C^\gamma f_{\beta\gamma}{}^\alpha = 0. \tag{4.36}$$

The second term vanishes because $f_{\beta\gamma}{}^\alpha$ is antisymmetric in β, γ, whereas $C^\beta C^\gamma$ is symmetric, hence what is left is

$$V_M\, \mathcal{C}_N{}^P V_P\, C^\alpha = 0. \tag{4.37}$$

Supposing that V_M and C^α are nonzero, this implies

$$\mathcal{C}_N{}^P V_P = 0, \tag{4.38}$$

that is,

$$\mathcal{C} V = 0. \tag{4.39}$$

So V lies in the kernel of $\mathcal{C}$. This is still not enough to prove triviality, though. Now, for the linear constraint:

$$X_{(MNP)} = 0. \tag{4.40}$$

Since $X_{MNP} = V_M \mathcal{C}_{NP}$, writing all index permutations this becomes

$$V_M \mathcal{C}_{NP} + V_N \mathcal{C}_{PM} + V_P \mathcal{C}_{MN} = 0. \tag{4.41}$$

Because this equation is tensorial, we may choose a basis adapted to V_M, assumed nonzero. Choosing to work in the canonical electric frame, for the rank-1 case I can choose V_M coordinates such that

$$V_M = v\,\delta_{M1}, \qquad v \in \mathbb{R} \setminus \{0\}. \tag{4.42}$$

That is, V_M is the vector with only the first component equal to v, while all other 55 components are zero. The next step is analysing the linear constraint component by component:

Case 1: $M = 1$, $N = i$, $P = j$ with $i, j > 1$. Then

$$0 = X_{(1ij)} = v\,\mathcal{C}_{ij}, \tag{4.43}$$

therefore

$$\mathcal{C}_{ij} = 0 \qquad (i, j > 1). \tag{4.44}$$

Case 2: $M = N = 1$, $P = j$ with $j > 1$. Then

$$0 = X_{(11j)} = v\,\mathcal{C}_{1j} + v\,\mathcal{C}_{j1} + 0. \tag{4.45}$$

Using symmetry $\mathcal{C}_{1j} = \mathcal{C}_{j1}$,

$$0 = 2v\,\mathcal{C}_{1j} \Longrightarrow \mathcal{C}_{1j} = 0 \qquad (j > 1). \tag{4.46}$$

Case 3: $M = N = P = 1$. Then

$$0 = X_{(111)} = 3v\,\mathcal{C}_{11} \Longrightarrow \mathcal{C}_{11} = 0. \tag{4.47}$$

Combining the three cases, we conclude

$$\mathcal{C}_{MN} = 0. \tag{4.48}$$

So the linear constraint kills all components of $\mathcal{C}_{MN}$. Since the symplectic form Ω_{MN} is nondegenerate,

$$\mathcal{C}_{MN} = 0 \quad \Longrightarrow \quad \mathcal{C}_N{}^P = 0. \tag{4.49}$$

Thus

$$C^\alpha (t_\alpha)_N{}^P = 0. \tag{4.50}$$

The fundamental **56** representation of $\mathfrak{e}_{7(7)}$ is faithful, which means that a linear combination of generators acts trivially on the **56** if and only if that linear combination is the zero element of the algebra. Therefore (4.50) implies

$$C^\alpha = 0. \tag{4.51}$$

The conclusion is then that the embedding tensor is null and the gauging is trivial:

$$\Theta_M{}^\alpha = V_M C^\alpha = 0. \tag{4.52}$$

□

4.4.2 The rank-2 case: reduction to a constrained form

The rank-2 case can be partially handled analytically as well. This case, however, serves also as an example of how fast the complexity of this analysis grows, since already for rank-2 the conditions to handle become more complex; a final computational check is needed to complete the analysis. The embedding tensor now takes the form:

$$\Theta_M{}^\alpha = V_M^{(1)} C_1^\alpha + V_M^{(2)} C_2^\alpha, \tag{4.53}$$

where $V_M^{(1)}$ and $V_M^{(2)}$ are linearly independent vectors in the fundamental **56** representation, and C_1^α, C_2^α are linearly independent adjoint vectors. Substituting this ansatz in the locality constraint gives:

$$\begin{aligned} 0 = (V_M^{(1)} C_1^\alpha + V_M^{(2)} C_2^\alpha)(V_N^{(1)} C_1^\beta + V_N^{(2)} C_2^\beta)\Omega^{MN} = V_M^{(1)} V_N^{(1)} C_1^\alpha C_1^\beta \Omega^{MN} + \\ + V_M^{(2)} V_N^{(2)} C_2^\alpha C_2^\beta \Omega^{MN} + V_M^{(1)} V_N^{(2)} C_1^\alpha C_2^\beta \Omega^{MN} + V_M^{(2)} V_N^{(1)} C_2^\alpha C_1^\beta \Omega^{MN}. \end{aligned} \tag{4.54}$$

The first two terms vanish identically, because $V_M^{(a)} V_N^{(a)}$, $a = 1, 2$, is symmetric in (M, N) while Ω^{MN} is antisymmetric. Hence

$$(C_1^\alpha C_2^\beta - C_2^\alpha C_1^\beta) V_M^{(1)} V_N^{(2)} \Omega^{MN} = 0. \tag{4.55}$$

Since C_1^α and C_2^α are linearly independent, otherwise this would again be the rank-1 case, the antisymmetrized coefficient does not vanish identically. Therefore

$$V_M^{(1)} V_N^{(2)} \Omega^{MN} = 0. \tag{4.56}$$

Thus the two-dimensional subspace spanned by $V^{(1)}$ and $V^{(2)}$ is isotropic. This is an example of how the isotropy conditions for the vectors emerge. One may therefore choose a Darboux basis in the canonical electric frame, such that

$$V_M^{(1)} = \delta_M{}^1, \qquad V_M^{(2)} = \delta_M{}^2. \tag{4.57}$$

Analogously to the previous case, this means that the two vectors have all null entries except for the first and the second one, respectively, that are each equal to one. Now, let us move onto the linear constraint, which gives useful simplifications. Defining

$$C_{1\,MN} := C_1^\alpha (t_\alpha)_{MN}, \qquad C_{2\,MN} := C_2^\alpha (t_\alpha)_{MN}, \tag{4.58}$$

where both $C_{1\,MN}$ and $C_{2\,MN}$ are symmetric tensors due to (4.32). With the basis choice (4.57), the tensor X_{MNP} becomes

$$X_{MNP} = \Theta_M{}^\alpha (t_\alpha)_{NP} = \delta_M{}^1 C_{1\,NP} + \delta_M{}^2 C_{2\,NP}. \tag{4.59}$$

The linear constraint then reads

$$\delta_M{}^1 C_{1\,NP} + \delta_M{}^2 C_{2\,NP} + \delta_N{}^1 C_{1\,PM} + \delta_N{}^2 C_{2\,PM} + \delta_P{}^1 C_{1\,MN} + \delta_P{}^2 C_{2\,MN} = 0. \tag{4.60}$$

Let indices i, j run over $\{3, \ldots, 56\}$. By analyzing components of (4.60), one obtains the following equations:

$$C_{1\,ij} = 0, \qquad C_{2\,ij} = 0, \qquad (i, j > 2), \tag{4.61}$$

$$C_{1\,1i} = 0, \qquad C_{2\,2i} = 0, \qquad (i > 2), \tag{4.62}$$

$$C_{1\,2i} + C_{2\,1i} = 0, \qquad (i > 2), \tag{4.63}$$

$$2C_{1\,12} + C_{2\,11} = 0, \tag{4.64}$$

$$C_{1\,22} + 2C_{2\,12} = 0, \tag{4.65}$$

$$C_{1\,11} = 0, \qquad C_{2\,22} = 0. \tag{4.66}$$

In matrix form, the most general solution of the linear constraint takes the form

$$C_{1\,MN} = \begin{pmatrix} 0 & a & 0 \\ a & b & c_i \\ 0 & c_i & 0 \end{pmatrix}, \qquad C_{2\,MN} = \begin{pmatrix} -2a & -\frac{b}{2} & -c_i \\ -\frac{b}{2} & 0 & 0 \\ -c_i & 0 & 0 \end{pmatrix}, \tag{4.67}$$

where $a = C_{112} = C_{121}, b = C_{122}$ and $c_i = C_{1i2} = C_{12i} = -C_{21i} = -C_{2i1}$ are real and they respectively belong to each piece of the block decomposition:

$$\{1\} \oplus \{2\} \oplus \{3, \ldots, 56\}.$$

Imposing the quadratic constraint (2.47) and defining the Lie algebra elements

$$\mathcal{C}_1 := C_1^\alpha t_\alpha, \qquad \mathcal{C}_2 := C_2^\alpha t_\alpha. \tag{4.68}$$

Then

$$(\mathcal{C}_a)_N{}^P = C_a^\alpha (t_\alpha)_N{}^P. \tag{4.69}$$

one gets the two equations

$$(\mathcal{C}_1)_N{}^1 C_1^\alpha + (\mathcal{C}_1)_N{}^2 C_2^\alpha + \delta_N{}^2 [C_1, C_2]^\alpha = 0, \tag{4.70}$$

$$(\mathcal{C}_2)_N{}^1 C_1^\alpha + (\mathcal{C}_2)_N{}^2 C_2^\alpha - \delta_N{}^1 [C_1, C_2]^\alpha = 0, \tag{4.71}$$

where

$$[C_1, C_2]^\alpha := C_1^\beta C_2^\gamma f_{\beta\gamma}{}^\alpha. \tag{4.72}$$

For $N = i > 2$, the Kronecker deltas vanish, and by linear independence of C_1^α and C_2^α one obtains

$$(\mathcal{C}_1)_i{}^1 = (\mathcal{C}_1)_i{}^2 = 0, \qquad (\mathcal{C}_2)_i{}^1 = (\mathcal{C}_2)_i{}^2 = 0. \tag{4.73}$$

To see other characteristics of the rank-2 case, one can use the relations presented when studying the electric frame in Section 2.5 [4]. For the embedding tensor of rank $n = 2$, the gauge indices

A, B, C can take only two values, either 1 or 2. The structure constants ${f_{AB}}^C$ are antisymmetric in the lower pair, ${f_{AB}}^C = -{f_{BA}}^C$, so the only potentially nonzero independent components are ${f_{12}}^1$ and ${f_{12}}^2$. Imposing the unimodularity condition ${f_{AB}}^B = 0$ for each value of A gives:

$$\begin{aligned} A = 1: &\quad {f_{11}}^1 + {f_{12}}^2 = {f_{12}}^2 = 0\,, \\ A = 2: &\quad {f_{21}}^1 + {f_{22}}^2 = -{f_{12}}^1 = 0\,. \end{aligned}$$

Hence ${f_{AB}}^C = 0$ for all $A, B, C \in \{1, 2\}$, and the rank-2 gauge algebra is necessarily **abelian**. Now, to analyze the form of the other tensors that make up X_{AN}^P in (2.60), I will apply the conditions (2.61) and (2.62), but first it is appropriate to fix notation: the index split is

$$A, B, C \in \{1, 2\}, \qquad c \in \{3, \ldots, 28\},$$

and magnetic duals are denoted by bars:

$$\bar{A} \equiv V_A, \quad \bar{a} \equiv V_a.$$

The 56-dimensional symplectic basis is ordered as

$$(V^1, V^2 \mid V^3, \ldots, V^{28} \mid V_1, V_2 \mid V_3, \ldots, V_{28}),$$

We introduce the shorthand vectors in $\mathbb{R}^{26}$:

$$\mathbf{h} := (h_{12}^3, \ldots, h_{12}^{28})^T, \qquad \mathbf{C} := (C_{12,3}, \ldots, C_{12,28})^T,$$

From equations (2.61), we obtain the following simplifications:

$${h_{(AB)}}^c = 0 \implies {h_{11}}^c = {h_{22}}^c = 0, \quad {h_{21}}^c = -{h_{12}}^c.$$

$$C_{(AB)c} = 0 \implies C_{11c} = C_{22c} = 0, \quad C_{21c} = -C_{12c}.$$

$$C_{(ABC)} = 0, \quad C_{A[BC]} = 0 \implies C_{111} = C_{222} = 0,$$

$$C_{121} = C_{112} = a, \quad C_{211} = -2a, \quad C_{212} = C_{221} = -\tfrac{b}{2}, \quad C_{122} = b.$$

The linear constraints $X_{A[MN]} = 0$, $X_{(AMN)} = 0$ fix all remaining components in terms of $a, b, \mathbf{h}, \mathbf{C}$:

$$X_{112} = X_{121} = a, \quad X_{122} = b, \quad X_{211} = -2a, \quad X_{212} = X_{221} = -\tfrac{b}{2},$$

$$X_{12c} = X_{1c2} = C_{12c}, \quad X_{12\bar{c}} = X_{1\bar{c}2} = h_{12}^c,$$

$$X_{21c} = X_{2c1} = -C_{12c}, \quad X_{21\bar{c}} = X_{2\bar{c}1} = -h_{12}^c.$$

Furthermore, with $f = 0$, the third equation of (2.62) reduces to

$$\begin{aligned} {f_{AB}}^E C_{ECD} - 4\, f_{(C[A}{}^E C_{B]D)E} + 4\, h_{(C[A}{}^a C_{B]D)a} &= 0, \\ \implies {h_{12}}^c\, C_{12c} &= 0, \end{aligned} \tag{4.74}$$

so $\mathbf{h} \cdot \mathbf{C} = 0$. The conclusions reachable analytically stop here; to conclude that no admissible gaugings exist at all for rank-2, a computational check is needed. For now, the only conclusions are that the gaugings have to be abelian, that the matrices C_{1MN} and C_{2MN} are very sparse, and the final condition (4.74).

4.5 Stochastic scan: code pipeline and realization

The first algorithm implemented to conduct the search for n-dimensional gauge subalgebras of $\mathfrak{e}_{7(7)}$ in $\mathcal{N} = 8$ supergravity is the **stochastic scan**. It has been realized entirely in Wolfram Mathematica. In this section, the main functions used for the code, as well as its computational limitations, will be presented step by step. In Figure 4.1, the whole pipeline for `rankn_gauging_analysis.m`, with entry point `RunRankN[n]`, is shown schematically. The pipeline has seven phases, each implemented by specific functions.

Phase 0: Loading the $\mathrm{E}_{7(7)}$ Data (`LoadRankNEnvironment`)

The first step consists in loading the source files for the generators and structure constants for $\mathrm{E}_{7(7)}$. `LoadRankNEnvironment["LoadLieART"→bool]` calls two sub-functions:

`LoadE7Data[]` , which is always called, reads from disk the precomputed $\mathrm{E}_{7(7)}$ data:

- `t`: a $133 \times 56 \times 56$ array of generator matrices in the fundamental representation, satisfying $\mathrm{Tr}(t_i t_j^T) = \delta_{ij}$.

  Note on the frame choice. Considering the set of generators $(t_\alpha)_M{}^N$ used, they are 56×56 matrices that encode the duality group $\mathrm{SL}(8, \mathbb{R})$, taken from [19], and it is one of the most known and convenient maximal subgroup of $\mathrm{E}_{7(7)}$ to employ. In this frame, the splitting $\mathbf{56} \to \mathbf{28} + \mathbf{28}'$ is preserved in block-diagonal form, and electric-magnetic mixing is forbidden. The adjoint representation of the exceptional group splits in the $\mathrm{SL}(8, \mathbb{R})$ as $\mathbf{133} \to \mathbf{63} + \mathbf{70}$, where $(t_{\mathbf{63}})_M{}^N$ are the block-diagonal generators that preserve this no-mixing condition, while the $(t_{\mathbf{70}})_M{}^N$ are off-diagonal and act as mixing operators. On the embedding tensor, this means that only the 63 electric components will actually be turned on, going from a 56×133 object to a 28×63 one.

- `structe7`: the $133 \times 133 \times 133$ structure-constant tensor $f_{\alpha\beta}{}^\gamma$.

- `metre7`: the 133×133 $\mathrm{E}_{7(7)}$ Cartan–Killing metric.

It also pre-computes the auxiliary array $(\Omega t_\alpha)_{MN} = \Omega_{MQ}(t_\alpha)^Q{}_N$ for all α, stored as `OmTa`, and verifies that the generators satisfy the symplectic condition ($|\, t_i\Omega + \Omega t_i^T| = 0$). This is a pure sanity check to verify that the generators in the source file are correct.

`LoadLieARTData[]` loads only if the boolean in the function is "True". It loads the `LieART` package, and pre-builds branching tables up to dimension 56 for the classical and exceptional subalgebras:

- $A_k = \mathfrak{sl}(k+1, \mathbb{C})$ for $k = 1, 2, 3, 4$,

- $B_n = \mathfrak{so}(2n+1, \mathbb{C})$ for $n = 2, 3$,

- $C_3 = \mathfrak{sp}(3, \mathbb{C})$,

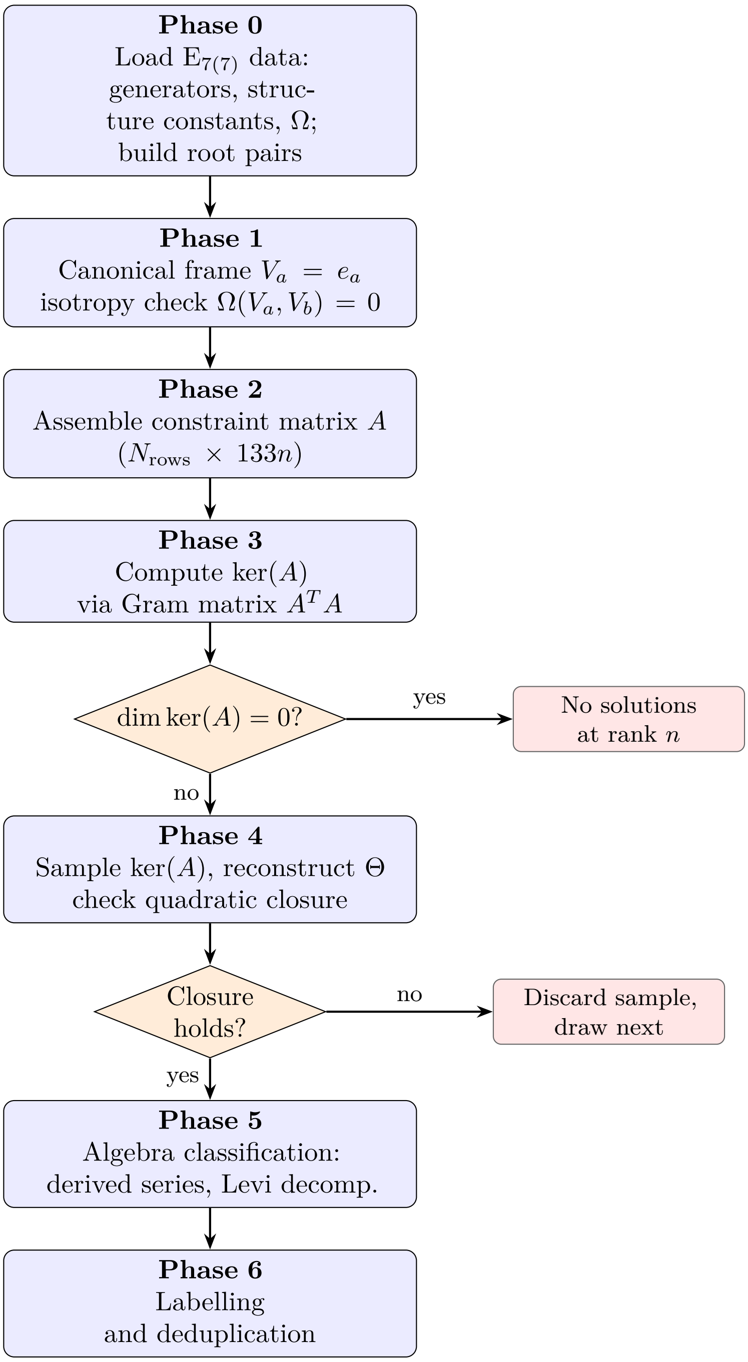

Figure 4.1: Pipeline of the RunRankN(n) algorithm for classifying rank-n gaugings in $\mathcal{N} = 8$ supergravity. The two decision nodes represent the main filters: the linear constraint (Phase 3) eliminates ranks for which no supersymmetric gauging exists, while the quadratic closure check (Phase 4) selects only consistent gauge algebras among the solutions of the linear system.

- $D_4 = \mathfrak{so}(8, \mathbb{C})$,
- G_2 (exceptional),

These tables are used in Phase 6 to match Casimir eigenvalues to named representations.

The two data objects are merged into a single `Association` object `env` that is passed through all subsequent phases.

`BuildRootPairData[e7Data, lieartData]` additionally scans `metre7` for off-diagonal non-zero entries, identifying the 63 step-generator pairs $(E_\alpha, E_{-\alpha})$ and their Dynkin weight vectors (read from the diagonal action of the Cartan generators via $f_{\alpha\beta}{}^{\beta}$). These are used in Phase 6 for orbit labelling.

Phase 1: Isotropy and the Canonical (`FixedRankNBasisData`)

The symplectic form Ω used is (2.15). In the canonical electric frame, $\Omega(e_i, e_j) = 0$ for all $i, j \leq 28$, since $V_a = e_a$ ($a = 1, \ldots, n$, standard unit vectors) automatically satisfies isotropy. **`CheckFixedRankNBasisCompatibility[env, n]`** performs a comprehensive structural check that includes (1) the isotropy condition via **`CheckIsotropyCondition[env, vectors]`**: it computes all pairwise symplectic products $V_a \cdot \Omega \cdot V_b$ and checks that each satisfies the isotropy constraint $|V_a \cdot \Omega \cdot V_b| <$ tol (tolerance default: 10^{-12}), (2) verification that the $\mathrm{E}_{7(7)}$ Cartan–Killing metric has correct dimensions 133×133, (3) confirmation that all basis vectors lie in the electric block (indices ≤ 28). This check ensures the entire computational framework is structurally sound before Phase 2 constructs the constraint matrix.

Phase 2: The Linear Constraint System (`BuildExactNRankSystem`)

Substituting the rank-n ansatz into the $X_{(MNP)} = 0$ constraint and contracting with a test triple $(v, w, u) \in (\mathbb{R}^{56})^3$ gives the linear equation (4.16) in the $133n$ unknowns $(C_1{}^{\alpha}, \ldots, C_n{}^{\alpha})$. Evaluating this with all r test triples produces a linear system of r equations. The full constraint matrix $\mathbf{A}$ has size $r \times 133n$, where r is the number of test triples used. For the canonical unit-vector frame, since each $V_a = e_a$ has support $\{a\}$, the support union has size n. For each support index $s \in \{1, \ldots, n\}$, the code enumerates all sorted triples (m, n, p) containing s as at least one element. After deduplicating the repeating triples, this yields all sorted triples touching the support $\{1, \ldots, n\}$, whose exact count is $\binom{58}{3} - \binom{58-n}{3} \approx O(n \cdot 56^2)$. All such triples are **enumerated exactly**, producing a dense but compact matrix. Test vectors are taken to be standard basis vectors, $v = e_i$, $w = e_j$, $u = e_k$, where $i, j, k \in \{1, \ldots, 56\}$, so that each row picks out one component of the constraint tensor:

$$e_i^M \, e_j^N \, e_k^P \, X_{(MNP)} \;=\; X_{(ijk)}. \tag{4.75}$$

If none of $\{M, N, P\}$ belongs to $\{1, \ldots, n\}$, all the Kronecker deltas vanish and $X_{(MNP)} = 0$ identically, regardless of the unknowns $C_a{}^{\alpha}$: those components bring no contribution and can therefore be discarded. The only non-trivial rows arise from triples (i, j, k) for which at least one index lies in the support $\{1, \ldots, n\}$ of the frame vectors; these are exactly the enumerated triples. After construction, however, as a final safeguard, the code removes any row that is numerically zero. The resulting matrix A therefore encodes every independent non-trivial component of $X_{(MNP)} = 0$ in the canonical frame, making the null-space computation exact and complete within this frame.

Phase 3: Finding the Null Space

To find the null space, the code computes the **Gram matrix** $G = A^T A$, which is significantly cheaper in terms of memory expense than a direct SVD (Singular Value Decomposition) computation. This is perfectly valid because $\ker(A) = \ker(A^T A)$: if $A^T A v = 0$ then $\|Av\|^2 = v^T A^T A v = 0$. The function `BuildExactNRankSystem` returns an `Association` that first calls `MatrixRank[mat, Tolerance` $\rightarrow 10^{-10}$`]` to compute the rank r_k of the constraint matrix. If one recovers full column rank, the function immediately returns `"NullDimension"` $\rightarrow$ `0` without any further computation. Otherwise, it computes the null space via the normal equations:

```
gramMat = Transpose[mat] . mat;
NullSpace[gramMat, Tolerance -> matTol]
```

It returns the null basis vectors as `"NullSpace"` and their count `Length[ns]` as `"NullDimension"`. If `NullDimension` $= 0$, this means there is no solution to the linear system of equations (4.16), and no non-trivial embedding tensor of rank n satisfying $X_{(MNP)} = 0$ in the canonical electric frame. The linear system has only trivial solutions and the scan immediately terminates for that rank.

Phase 4: Closure Optimisation (`SampleNullSpaceClosureN`)

A generic point in the null space of A satisfies the linear constraint (2.49) but does not necessarily close as a Lie algebra. The next step is therefore a check for closure. Given a null-space basis $\{N_i\} \subset \mathbb{R}^{133n}$, a point $\lambda = (\lambda_1, \ldots, \lambda_d)$ gives generator coefficient vectors

$$\vec{C} = \sum_{i=1}^{d} \lambda_i N_i \;\in\; \mathbb{R}^{133n}, \qquad C_a{}^{\alpha} = (\vec{C})_{(a-1)\cdot 133+\alpha}. \tag{4.76}$$

The gauge generators are $G_a = C_a{}^{\alpha}\, t_\alpha$ discussed in Section 4.3. The closure condition $[G_a, G_b] = f_{ab}{}^{c} G_c$ is quadratic in the null-space coordinates $\lambda \in \mathbb{R}^d$. In the code, the actual commutator used for the closure check keeps its original form:

$$[G_a, G_b] = C_a{}^{\alpha} C_b{}^{\beta} [t_\alpha, t_\beta] = C_a{}^{\alpha} C_b{}^{\beta} f_{\alpha\beta}{}^{\gamma}\, t_\gamma. \tag{4.77}$$

This is never evaluated as a product of 56×56 matrices; instead, `fSFlat` is a pre-computed $(133^2 \times 133)$ reshape of the structure-constant tensor, so `fSFlat . Transpose[genCoeffs]` yields all n commutator vectors $[G_a, G_b]$ directly in the 133-dimensional adjoint basis in a single BLAS (Basic Linear Algebra Subprograms, a standard library of highly optimized linear algebra routines) call, without touching the fundamental-representation matrices.

Let $U \in \mathbb{R}^{n\times 133}$ be the matrix whose rows form an orthonormal basis of $\operatorname{span}\{C_a\}$ in $\mathbb{R}^{133}$, the coordinate representation of $\operatorname{span}\{G_a\} \subset \mathfrak{e}_{7(7)}$ in the basis $\{t_\alpha\}$, computed via `Orthogonalize`. If the algebra closes, the commutator must lie in $\operatorname{span}\{G_a\}$. If the algebra closes, every commutator must lie in $\operatorname{span}\{G_a\}$. The generator stack is first normalised globally ($C \to C/\|C\|$) so $\sum_a \|G_a\|^2 = 1$, which fixes the overall scale of λ, therefore making the closure check scale-independent. To check for which λ the algebra closes I employ the **closure residual**:

$$F(\lambda) = \sum_{a,b=1}^{n} \left\| [G_a, G_b]_{\perp,\lambda} \right\|^2, \tag{4.78}$$

where $[G_a, G_b]_{\perp,\lambda}$ denotes the component of the commutator vector in the adjoint basis, orthogonal to $\operatorname{span}\{G_a\}$: $[G_a, G_b]_{\perp,\lambda} = v_{ab}(\lambda) - U^T(U v_{ab}(\lambda))$ with $v^{\gamma}_{ab}(\lambda) = C_a{}^{\alpha}(\lambda) C_b{}^{\beta}(\lambda) f_{\alpha\beta}{}^{\gamma}$. The closure residual satisfies $F(\lambda) = 0$ if and only if the algebra closes. Using the Gram–Schmidt decomposition rather than the full pseudoinverse $f_{ab}{}^{c} = \left(C^T\right)^{+}{}_{c}{}^{\gamma}\left(C_a{}^{\alpha} C_b{}^{\beta} f_{\alpha\beta}{}^{\gamma}\right)$, costs $O(133n^2)$ per evaluation. This replaces the $O(133^2 n)$ SVD that a full pseudoinverse would require at every step when looking for solutions via the BFGS (Broyden–Fletcher–Goldfarb–Shanno) algorithm, explained in the next paragraph. This choice makes each evaluation significantly cheaper. A special case in the computation happens when the null space is one-dimensional ($d = 1$), and the solution is unique up to sign and no optimisation is needed. The code takes a dedicated fast path: it evaluates F at

$\lambda = \{1\}$ directly and skips the entire BFGS machinery.

BFGS hybrid strategy. The closure residual $F(\lambda)$ is a smooth function of the d null-space coordinates, and finding its zeros is a nonlinear optimization problem. The code uses **BFGS**, a quasi-Newton method that minimises a smooth function by iteratively building an approximate Hessian from successive gradient differences, without ever computing it explicitly. Mathematica's `FindMinimum` with `Method` $\to$ `"QuasiNewton"` implements standard BFGS, which stores and updates a dense approximate Hessian over the d-dimensional parameter space, making it well-suited for null spaces of moderate dimension. Since F can have many local minima, especially in high-dimensional null spaces, a single BFGS run from a random starting point may converge to a non-closure basin. `SampleNullSpaceClosureN` therefore runs a **multi-start hybrid** strategy:

1. **Pre-screening.** A large pool of random unit vectors $\lambda \in S^{d-1}$ is drawn and F is evaluated at each without any optimisation.
2. **Adaptive seeds.** The best N_{adapt} pre-screened points, specifically those with the lowest F before any optimisation, are used as warm starts for `FindMinimum` (`MaxIterations` $\to$ `200`, `AccuracyGoal` $\to$ `8`).
3. **Purely random complement.** The remaining $N_{\text{starts}} - N_{\text{adapt}}$ starts use fresh random seeds that bypass the pre-screen ranking entirely. This is an essential step, because pre-screening can miss narrow basins of attraction that happen to score poorly prior to optimization but that are genuine solutions.

Each start has a time cap implemented via `TimeConstrained`, with a time budget that increases with the dimension of the null space. When a local minimum with $F < 10^{-3}$ is found, the structure constants are extracted. Since closure holds, the commutator vector $v_{ab}^{\gamma} = C_a{}^{\alpha} C_b{}^{\beta} f_{\alpha\beta}{}^{\gamma}$ lies in $\mathrm{span}\{C_c{}^{\gamma}\}$, so the system $f_{ab}{}^{c} C_c{}^{\gamma} = v_{ab}^{\gamma}$ is consistent. The Moore–Penrose pseudoinverse $(C^T)^+$ gives the exact solution to (4.25), computed via a single `PseudoInverse` call (SVD) per confirmed solution. This call is deliberately deferred to the end, only for confirmed closed candidates, for efficiency.

Phase 5: Algebra [20](`algebraInvariantDataN`, `classifyNDAlgebraType`)

From the $n \times n \times n$ structure-constant tensor $f_{ab}{}^{c}$, several invariants are computed. Here they are all reported and their definition can be found in Appendix **??**.

Adjoint representation and series. The adjoint matrices $(\mathrm{ad}_{G_a})_{cb} = f_{ab}{}^{c}$ are built by `adjointMatricesFromStructureN`.

- **Derived dimension** (`DerivedDim`): $\dim[\mathfrak{g}, \mathfrak{g}]$, the rank of the span of all brackets $[G_a, G_b]$.
- **Nilpotency**: the lower central series $\mathfrak{g}^{k+1} = [\mathfrak{g}, \mathfrak{g}^k]$ is iterated until the dimension stabilises. `NilpotentQ` $=$ `True` if the series reaches 0.
- **Solvability**: the derived series $\mathfrak{g}^{(k+1)} = [\mathfrak{g}^{(k)}, \mathfrak{g}^{(k)}]$ is iterated. `SolvableQ` $=$ `True` if it reaches 0.
- **Unimodularity**: `Unimodular` $=$ `True` if $\mathrm{tr}(\mathrm{ad}_{G_a}) = f_{ab}{}^{b} = 0$ for all a.

Killing form.

$$K_{ab} = \mathrm{tr}(\mathrm{ad}_{G_a} \circ \mathrm{ad}_{G_b}) = f_{ac}{}^{d}\, f_{bd}{}^{c}. \tag{4.79}$$

The eigenvalues of K are computed as well, and the triple (p, q, z) counts positive, negative, and zero eigenvalues at threshold 10^{-8}. `SemisimpleQ` = `True` if $z = 0$, meaning that the Cartan–Killing form is non-degenerate, and by Cartan's criterion this implies semisimplicity.

$\mathrm{E}_{7(7)}$ **Gram signature.** The $n \times n$ Gram matrix in the $\mathrm{E}_{7(7)}$ metric, $G_{ab} = C_a{}^{\alpha}\, (\text{metre7})_{\alpha\beta}\, C_b{}^{\beta}$, has signature `E7GramSignature`= (p', q', z'). This distinguishes algebras with the same Killing signature but different embeddings in $\mathfrak{e}_{7(7)}$.

Dynkin index and Casimir eigenvalues. The generator matrices $G_a = C_a{}^{\alpha} t_{\alpha}$ (size 56×56) are assembled and the quadratic Casimir in the **56** is computed:

$$C_2^{(\mathbf{56})} = \sum_{a=1}^{n} G_a^2. \tag{4.80}$$

Its trace gives the **Dynkin index**

$$I = \frac{\mathrm{tr}_{\mathbf{56}}(C_2^{(\mathbf{56})})}{\mathrm{tr}_{\mathbf{56}}(\sum_{\alpha=1}^{133} t_{\alpha}^2)}, \tag{4.81}$$

stored as `DynkinIndexRational` = `Rationalize[I, 0.001]`. The 56 eigenvalues of $C_2^{(\mathbf{56})}$ are stored in `Branching56Evals`.

Algebra type labelling. `classifyNDAlgebraType[invData, n, fABC]` assembles the string label `AlgebraType` from the invariants alone, without using LieART: for semisimple algebras it uses the Killing signature and Dynkin index to produce a label such as `"so(5,3)"`; for non-semisimple algebras it uses the Levi structure and Killing signature to signal a Levi decomposition.

Phase 6: Labelling and Deduplication (`ClassifyWithLieARTN`, `DeduplicateRankNSolutions`)

After Phase 5 produces the algebra type and invariants, each confirmed solution is passed to `ClassifyWithLieARTN[genCoeffs, rootPairData, env]`. This function classifies every generator $G_a = C_a{}^{\alpha} t_{\alpha}$ individually by decomposing its coefficient vector into Cartan ($\alpha = 1, \ldots, 7$) and step-generator ($\alpha = 8, \ldots, 133$) parts:

- If the Cartan norm dominates ($> 90\%$ of total norm), the generator is labelled `Cartan:`k where k is the index of the dominant Cartan generator.
- If the step norm dominates, the generator is matched to the nearest root pair $(E_{\alpha}, E_{-\alpha})$ in `rootPairData` and labelled `E7root:[`$w_1, \ldots, w_7$`]` using the Dynkin weight vector.
- Otherwise it is labelled `Mixed`.

The result is stored as `LieARTData` in the solution association, providing an orbit-level fingerprint of the embedding inside $\mathfrak{e}_{7(7)}$ beyond what the Killing form and Dynkin index alone distinguish. The LieART embedding identification is restricted to semisimple algebras. Consequently, for Levi-decomposable solutions, returns no match is returned (`Branching56Named`= `"N/A"`).

Finally, `DeduplicateRankNSolutions[solutions, n]` groups all confirmed solutions by the composite key {`AlgebraType`, `E7GramSignature`}. The algebra type label separates most cases, while the $E_{7(7)}$ Gram signature catches residual ambiguities between algebras with identical Killing numbers but different embeddings in $\mathfrak{e}_{7(7)}$. One representative per equivalence class is retained in the final results.

4.5.1 Computational limits and error analysis

The pipeline combines exact linear algebra, used in the rank and null-space determination steps, with a stochastic nonlinear search employed during closure optimization. These two stages have very different reliability profiles, and it is important to state precisely what the code does and does not guarantee. This subsection collects the numerical tolerances, the propagation of floating-point error, and the structural limitations of the method.

All heavy linear algebra is performed in double precision ($\varepsilon_{\text{mach}} \approx 2.2 \times 10^{-16}$). There is a `"WorkingPrecision"` → `50` option as well, reserved for the few exact symbolic checks and does not govern the main floating-point kernels. The code uses a deliberately layered set of tolerances rather than a single global value, because the operations they control have different noise floors, shown in the Table 4.1.

Stage	Tolerance	Role
Linear system / rank	10^{-10}	SVD singular-value cutoff in `MatrixRank`/`NullSpace`
Null space (normal eqns.)	10^{-5}	effective σ-cutoff after squaring (see below)
Closure residual	10^{-3}	accept/reject threshold for $F(\lambda)$
Isotropy check	10^{-12}	$\max_{a,b} \lvert\Omega(V_a, V_b)\rvert$ at frame setup
LieART Casimir match	10^{-8}	eigenvalue grouping in branching identification

Table 4.1: Tolerance values for every step of the code.

The values are not arbitrary: each tolerance is chosen to sit at least one order of magnitude above the backward-error floor, which is the background noise created by the computer rounding off decimals. At the same time, tolerance is also far below the smallest physically meaningful gap, so that genuine and spurious quantities are cleanly separated.

For the computation of the null space dimension, the number of independent gauge candidates equals $d = \dim\ker A$, the null-space dimension of the $r \times 133n$ constraint matrix A. Because A is too large to feed directly to `NullSpace`, the code forms the Gram matrix $G = A^{\top}A$ and computes $\ker A = \ker G$ via eigendecomposition. This is mathematically exact, but it squares the singular values: a relative tolerance $\tau = 10^{-10}$ applied to G corresponds to a cutoff $\sigma < \sqrt{\tau}\,\sigma_{\max} = 10^{-5}\sigma_{\max}$ on A itself.

As for the acceptance threshold for the closure residual (4.78), it is 10^{-3}, which is much larger than the 10^{-8}–10^{-10} tolerances above. This is not a loss of rigour but a direct consequence of how $F(\lambda)$ is computed. The orthogonal component $[G_a, G_b]_{\perp} = v_{ab} - U^{\top}(Uv_{ab})$ is obtained by subtracting two nearly equal vectors, an operation subject to catastrophic cancellation. The resulting noise floor scales roughly as

$$F_{\min} \sim n\,(\dim\mathfrak{g})^{3/4}\,\sqrt{\varepsilon_{\text{mach}}} \qquad (4.82)$$

and for $n = 28$, for example, a genuinely closed algebra still registers $F \sim 10^{-5}$. Since non-closing candidates have $F \sim 0.01$–1, the threshold 10^{-3} separates the two populations by roughly two orders of magnitude in either direction. A tighter threshold would reject true solutions, while a looser one would admit near-misses. The code therefore sets `closureTol = Max[tol, 1e-3]`.

Algorithm limitations

The most important limitation of this method is that Phase 4 is a search that uses sampling, therefore it is not exhaustive, despite the attempts to mitigate this with the different BFGS approaches described. Nevertheless, a sufficiently small basin of attraction could in principle be missed. This means that, when no algebras are found but the dimension of the null space is different from zero, this does not necessarily mean that there are no algebras at all, they might simply be hard to find. These limits are exactly what motivates us to try also an exact scan. It might seem a hard task at first glance, since it requires exact solving of the closure system, but the particular $\mathrm{SL}(8,\mathbb{R})$ frame in which we have chosen to operate provides an especially convenient environment that makes this method absolutely feasible.

However, for the cases of ranks 1–7, as will be shown in the next chapter, this scan finds that the null space has always dimension zero, therefore the stochastic phase of the code is never entered and results are not influenced by these matters.

4.5.2 Computational Infrastructure: Cloud Virtual Machine

For low ranks (1 to 7), which do not require the null-space computation, the scans were run locally under the Mathematica desktop kernel. For cases 13 to 28, however, the combination of multiple factors such as a large null space ($d \gg 1$), a high-dimensional BFGS optimisation, and the need for hundreds of independent random starts resulted in a total wall time of 17–20 hours and made local execution impractical.

These runs were therefore executed on a Ubuntu virtual machine hosted on the **CloudVeneto** academic cloud infrastructure. The code was run with extended precision `WorkingPrecision → 50` and `NStartPoints → 400`, and a fresh random seed at each invocation.

4.6 Exact scan: motivation and methods

The companion implementation `exact_scan.wl` reproduces the linear stage of the BFGS pipeline verbatim, but replaces the stochastic closure search with an exact symbolic treatment. In this section, a presentation of this new approach will be done, as well as underlining the motivation for its employment.

The exact scan attempts to resolve the closure check symbolically. It has the additional aim to try to retrieve the entire family structure of the gaugings: the BFGS scan only finds certain points of continuous solutions and identifies them, but cannot see the whole family. The closure check (4.78) is computed symbolically as quadratic polynomials in the projective coordinates λ,

$$[G_a, G_b](\lambda) = v_{ab}{}^{\gamma}(\lambda)\, t_{\gamma}, \tag{4.83}$$

where $\gamma = 1, \ldots, 133$. Only a subset of these adjoint directions is ever populated by the generators

G_a of a given family. We define the support S as the set of directions γ for which some $G_a(\lambda)$ has a non-zero component (for example, at rank 28, $|S| = s = 63$ of the 133), while the remaining $133 - s$ directions vanish identically on the whole family. Closure is then verified in two complementary pieces. First, for every adjoint index $\gamma \notin S$ one checks the exact polynomial identity $v_{ab}{}^{\gamma}(\lambda) \equiv 0$. Of course, usually the commutator can have directions that lead outside the support, and verifying that this does not happen is a necessary condition for closure. This condition is checked to be satisfied identically for every rank 7–28. The second step verifies that, within the support, the bracket lies in the linear span of the n generators themselves. This check is necessary because the commutator might stay in the support but exit the span of the n generators. Let us remind that the generator's structure is:

$$G_a(\lambda) \;=\; \sum_{i=1}^{d} \lambda_i \,(\mathrm{N}_i)_a{}^{\gamma}\, t_\gamma, \qquad a = 1, \dots, n, \tag{4.84}$$

where the N_i span the null space $\ker(\mathbf{A})$. Inserting this parametrization into the closure constraint, one gets:

$$[\, G_a(\lambda),\, G_b(\lambda)\,] \;=\; \sum_{i,j=1}^{d} \lambda_i\lambda_j \,(\mathrm{N}_i)_a{}^{\gamma}(\mathrm{N}_j)_b{}^{\delta}\, f_{\gamma\delta}{}^{\varepsilon}\, t_\varepsilon, \tag{4.85}$$

Therefore, asking for closure means checking if there are values of λ that bring what lies in the complement of $\mathrm{span}\{G_1(\lambda), \dots, G_n(\lambda)\}$ to zero. Here $\Delta_{ab}{}^{\varepsilon}(\lambda)$ represents it component-wise:

$$[\, G_a(\lambda),\, G_b(\lambda)\,] \;=\; C_{ab}{}^{\varepsilon}(\lambda)\, t_\varepsilon, \qquad C_{ab}{}^{\varepsilon}(\lambda) \;=\; \sum_{i,j} \lambda_i\lambda_j \,(\mathrm{N}_i)_a{}^{\gamma}(\mathrm{N}_j)_b{}^{\delta}\, f_{\gamma\delta}{}^{\varepsilon}, \tag{4.86}$$

$$\Delta_{ab}{}^{\varepsilon}(\lambda) \;=\; [\Pi_\perp(\lambda)]^{\varepsilon}{}_{\zeta}\, C_{ab}{}^{\zeta}(\lambda) \;=\; [\Pi_\perp(\lambda)]^{\varepsilon}{}_{\zeta}\, ([\, G_a(\lambda),\, G_b(\lambda)\,])^{\zeta}, \tag{4.87}$$

where $\Pi_\perp(\lambda)$ is the λ-dependent projector. Closure happens if $\Delta_{ab}{}^{\varepsilon}(\lambda) = 0$. However, by splitting this factor into the piece inside of the support S of the generators and the one outside this condition is easier to manage. The out-of-support condition $\Delta_{ab}^{\varepsilon \notin S}(\lambda) = 0$ is empirically found to be satisfied by any λ, and this convenient cancellation brings relevant simplification to the problem. Looking at this condition directly, it is clear that it is automatically satisfied for any λ if the support is a subalgebra of $\mathfrak{e}_{7(7)}$, so that the structure constants $f_{\gamma\delta}{}^{\varepsilon} = 0$ whenever the ε index lies outside the support.

$$\Delta_{ab}^{\varepsilon \notin S}(\lambda) = \sum_{i,j=1}^{d} \lambda_i\lambda_j\, M_{(ia)(jb)}{}^{\varepsilon}, \qquad M_{(ia)(jb)}{}^{\varepsilon} = (\mathrm{N}_i)_a{}^{\gamma}(\mathrm{N}_j)_b{}^{\delta} f_{\gamma\delta}{}^{\varepsilon}, \tag{4.88}$$

What happens is that the support is verified to lie inside $\mathfrak{sl}(8)$, and since $[\mathfrak{sl}(8), \mathfrak{sl}(8)] \subseteq \mathfrak{sl}(8)$, this condition is automatically satisfied always. So now the closure condition rests entirely on the in-support part. To certify this, we conduct a rank test on the row-stacked matrix

$$M(\lambda) = \begin{bmatrix} G(\lambda) \\ B(\lambda) \end{bmatrix}, \qquad G(\lambda) \in \mathbb{Q}[\lambda]^{n\times s}, \quad B(\lambda) = \big([G_a, G_b](\lambda)\big)_{a<b} \in \mathbb{Q}[\lambda]^{\binom{n}{2}\times s}, \tag{4.89}$$

and the condition on λ becomes:

$$\delta(\lambda) := \mathrm{rank}\, M(\lambda) - \mathrm{rank}\, G(\lambda) = 0 \quad \Longleftrightarrow \quad \mathrm{rowspan}\, B(\lambda) \subseteq \mathrm{rowspan}\, G(\lambda), \tag{4.90}$$

Basically, if appending the brackets does not make the rank any higher, the algebra is certified to close, since no other generators are added to the rank count. Since the symbolic rank equals the largest size of a minor in the matrix that is a nonzero polynomial, the identity

$$\operatorname{rank}_{\mathbb{Q}(\lambda)} M = \operatorname{rank}_{\mathbb{Q}(\lambda)} G \ = r_G \tag{4.91}$$

is equivalent to the vanishing of all $(r_G{+}1) \times (r_G{+}1)$ minors of M as polynomials in λ. For all ranks $6-17$ the result $\delta = 0$ is certified for every λ. Therefore, the set of consistent gaugings is the full projective family $\mathbb{RP}^{d-1}$, and no nonlinear search is needed for these cases. The d basis vectors N_i are the family. Every consistent gauging of that rank is a point λ in that space, since $G_a(\lambda) = \sum_{i=1}^{d} \lambda_i\, (\mathrm{N}_i)_a{}^\gamma\, t_\gamma$. The family is represented by the span of these vectors. Overall rescaling $\lambda \to c\lambda$ simply scales the coupling, therefore giving the same gauging, so the genuine family is $\mathbb{RP}^{d-1}$, not $\mathbb{R}^d$. However, while the set of null vectors fixes the family, different λs can realize different specific algebras. Unfortunately, for ranks 18-28 we find $\delta \neq 0$. This means that, to find a family, one must solve the equations in λ precisely, but this is unfeasible for higher ranks. However, the $\delta = 0$ condition reveals that $\mathcal{Z} = \{\lambda : \delta(\lambda) = 0\}$, viewed in the parameter space $\mathbb{R}^d$ of the coefficients $\lambda = (\lambda_1, \ldots, \lambda_d)$, is not a generic algebraic hypersurface but a finite union of coordinate subspaces. Writing $\{e_1, \ldots, e_d\}$ for the standard basis of this space, where the axis e_p activates the null vector N_p,

$$\mathcal{Z} \ = \bigcup_S \operatorname{span}\{e_p : p \in S\}, \tag{4.92}$$

the union running over index sets S. Concretely, whenever a single generic point of $\operatorname{span}\{e_p : p \in S\}$ closes, every point of that subspace closes (this was checked by verifying that $\delta \equiv 0$ on the flat on many random integer combinations; it is a heuristic verification but still robust, and gives coherent results). Hence the problem of mapping out $\mathcal{Z}$ reduces to the finite, combinatorial task of finding the maximal closing index sets S. We enumerate these by a greedy growth procedure: starting from each closing singleton we repeatedly adjoin a coordinate that preserves $\delta = 0$, over several randomized orders, and retain only the maximal sets.

At a real point λ we extract a maximal set of independent rows of $\mathrm{G}_a(\lambda)$ as a basis of the gauge algebra, and recover its own structure constants, then we build the gauge algebra's own Killing form. Its eigenvalue signature (p, q, z) fixes the labels. Sweeping the sign sectors of the free parameters λ_p $(p \in S)$ enumerates all real forms a family realises.

The exact scan was introduced precisely to overcome the limitation exposed in Section 5.2.2: in this new scan, closure is upgraded from a set of isolated solutions found by sampling to a polynomial identity valid on the whole family, removing the residual ambiguity of the $\sim 10^{-4}$ BFGS closure failures at rank 28. However, this method seems to be convenient only when the underlying structure of the frame used presents particular symmetries, that might impose simplifications on the non-linear and usually complex systems in λ.

Chapter 5

Results: no-go theorem and consistent gaugings

The results obtained from the apparatus described in Chapter 4 to conduct the rank-n scan are now presented systematically. The most fascinating result, presented as a **no-go theorem**, is the complete absence of gaugings for ranks 1 to 6, studied in the purely electric frame. Unlike the higher-rank results, which are obtained by a stochastic search and are therefore existence statements, the low-rank claim is a closed theorem: it follows from a single exact rank computation, and the only quantity that enters is the dimension of the null space of the linear constraint matrix. This section provides this result precisely, presenting its nuances and meaningfulness. To extend the generality of the theorem to all frames, a computational check of the minimum support allowed by the linear constraint (2.49) is added, giving very strong evidence that this result is valid universally. Apart from this last check, all implementation details for other results are deferred to the previous chapter.

5.0.1 No-go theorem

Theorem 5.1. Let ${\Theta_M}^\alpha \in \mathbf{56} \otimes \mathbf{133}$ be an embedding tensor of $\mathcal{N} = 8$ supergravity in 4d, satisfying the linear constraint $X_{(MNP)} = 0$ of Eq. (2.49) and the quadratic locality constraint of Eq. (2.48). Suppose Θ has rank n as a 56×133 matrix with $1 \leq n \leq 6$. This implies: $\Theta = 0$. Equivalently, there is no non-trivial gauge algebra of dimension $n \leq 6$ that can be realised in the canonical, purely electric frame of $\mathrm{E}_{7(7)}$.

The restriction to the electric frame retains a certain generality. The locality constraint (2.48) implies that the charge vectors span an isotropic subspace of $(\mathbb{R}^{56}, \Omega)$, and a symplectic rotation can always bring any such subspace into the electric block. Two embeddings related by an $\mathrm{E}_{7(7)}$ transformation define the same physical gauging, so it is sufficient to prove the theorem for the chosen canonical electric frame $V_a = e_a$ in $\mathrm{SL}(8, \mathbb{R})$ for the same conclusion to be valid also for any other frame in the same orbit of $E_{7(7)}$.

Proof. The main advantage of this proof is the possibility to transpose the question of whether there are any gaugings at all for $\mathrm{rank}(\Theta) \leq 6$, to a statement on the null space of a certain constraint matrix, exactly the one introduced in (4.17). Therefore, with the rank-n ansatz ${\Theta_M}^\alpha = \sum_{a=1}^{n} (V_a)_M \, {C_a}^\alpha$, in the canonical electric frame $V_a = e_a$, the linear constraint $X_{(MNP)} = 0$ becomes the homogeneous system

$$A\,\mathbf{C} = \mathbf{0}, \qquad A \in \mathbb{R}^{r\times 133n}, \qquad \mathbf{C} = ({C_1}^{\alpha}, \ldots, {C_n}^{\alpha}) \in \mathbb{R}^{133n}, \tag{5.1}$$

whose explicit rows are given in Eq. (4.16). A non-trivial embedding tensor of rank n exists if and only if this system admits a non-zero solution. From a linear algebra perspective, this happens if and only if

$$d \;:=\; \dim(\ker A) \;>\; 0. \tag{5.2}$$

Proving the statement for this theorem is therefore equivalent to the assertion that A has **full column rank** for every $n \in \{1, \ldots, 6\}$. To assure the soundness of this proof, one must first check that the rows assembled by the code represent every one of the independent components of the constraint, providing completeness to the argument. Then, we will also show how, due to the specific parts of the code involved, this theorem is numerically sound as well. The other relevant information about how the constraint matrix A is built, and how the rank of its kernel is found, is exactly that of Phases 2–3 of Chapter 4 and is not repeated here.

5.0.2 Completeness of triples enumeration

One thing that makes the choice of the canonical electric frame particularly convenient is that, due to the sparseness of the vectors $V_a = e_a$, $a = 1, ..., n$, it is possible to enumerate every row of the linear system $X_{(MNP)} = 0$, since even for the highest $\mathrm{rank}(\Theta) = 28$, the total number of rows is manageable comfortably, and there is no need to use just a limited batch of sampling rows, leaving some of them out. This is precisely what makes this an actual proof. The ansatz ${\Theta_M}^{\alpha} = \sum_a \delta_{aM} {C_a}^{\alpha}$ renders the symmetrised constraint gives

$$X_{(MNP)} \;=\; \tfrac{1}{3} \sum_{a=1}^{n} \Big[\delta_{aM}\, {C_a}^{\alpha} (t_\alpha)_{NP} + \delta_{aN}\, {C_a}^{\alpha} (t_\alpha)_{MP} + \delta_{aP}\, {C_a}^{\alpha} (t_\alpha)_{MN} \Big]. \tag{5.3}$$

The Kronecker deltas restrict every non-zero contribution to triples (M, N, P) for which at least one index lies in the support $\{1, \ldots, n\}$ of the frame, as already explained in the previous section. The number of test triples (and subsequently rows) to employ is given by

$$r(n) = \binom{58}{3} - \binom{58-n}{3} \tag{5.4}$$

counting all ordered triples (u, v, w) with $1 \le u \le v \le w \le 56$ that intersect $\{1, \ldots, n\}$. Running the pipeline for each $n \in \{1, \ldots, 6\}$ returns

$$d = \dim\ker A = 0, \tag{5.5}$$

in every case. The constraint matrix has **full column rank**, and the scan terminates before the closure phase. There is no rank-$n \le 6$ embedding tensor that satisfies $X_{(MNP)} = 0$ in the canonical electric frame. By the completeness argument above, this concludes the proof, provided that the method used to compute the rank is reliable.

This reliability was established in detail in Sec. 4.5.1, and only the conclusion is recalled here. The rank is read off from the singular values of A with a relative cutoff $\tau = 10^{-10}$, while the backward error of the double-precision computation is bounded by $\sim 10^{-13}\sigma_{\max}$ for the largest case ($n = 6$). The cutoff, therefore, sits roughly 600 times above the error floor, so no genuine non-zero singular value can be mistaken for zero. The built-in sensitivity guard, which recomputes the rank at the looser tolerance 10τ and flags any disagreement, produced no warning for any $n \leq 6$, confirming that all singular values of A lie at least four orders of magnitude above the threshold. The full-rank verdict is thus certified, with the only residual uncertainty being the accuracy of the stored generators themselves, however this is separately validated by the symplecticity check performed during the loading Phase 0 of the code and satisfied exactly. Under these conditions, the probability of a false full-rank declaration is negligible: it would require the near-zero singular values to exceed the threshold by more than ~ 600 times the backward-error bound. This proof can therefore be considered numerically sound. $\square$

5.1 Minimum support analysis

The generalization of the theorem over all frames of $E_{7(7)}$ is given by a computational calculation for the minimum support allowed by the linear constraint system $\mathbf{A}\Theta = 0$ that binds the embedding tensor to be in the 912 representation. For this verification, no assumption on Θ was done. The aim is to find the first rank for Θ that satisfies the linear condition, so the approach is the same as before but without choosing any specific frame for Θ. A first BFGS check estimates $r_{\min}$ and test for the presence of a prescribed rank by sampling $\Theta = \lambda \cdot \texttt{basis912}$ at random $\lambda \in \mathbb{R}^{912}$ and evaluating $\operatorname{rank}\Theta$, refining towards the minimising rank.

The more robust check, however, is the pipeline that continues on Python with LAPACK, for efficiency: it reads the sparse triplets of A and obtains an orthonormal kernel basis from the economy SVD $\Theta = U\Sigma V^{\mathsf{T}}$ (the right-singular vectors with $\sigma \leq 10^{-9}$ spanning $\ker A$), then performs an iterative pattern. Basically, the algorithm assumes Θ of a desired rank, via SVD, by putting to zero the singular value that increase its rank. By putting this matrix into the linear constraint and checking by how much it does not satisfy it, the code gives back another Θ that is no longer rank-deficient but the singular value(s) that made it so are smaller. By iterating these steps, if the singular values converge to zero, with a residual smaller than 10^{-5}, this means that that rank is the first one to admit gauging. Both tests show that the result is exactly what was expected: the lowest possible rank admissible is 7, and it converges definitely at machine precision. For lower gaugings, this residual is always at around 10^{-1} or higher.

5.2 Detailed Results

The threshold $n = 7$ at which gaugings first appear is a hard boundary of the electric frame, below which the linear constraint $X_{(MNP)} = 0$ alone forbids the existence of non-trivial embedding tensors. When this boundary is crossed, however, the null space of the constraint matrix starts having $\dim(\ker A) > 0$, all reported in Table 5.2, and the computation continues, moving to the closure check.

Table 5.1: Frame-independent minimum-rank test in the **912** of $E_{7(7)}$: smallest relative residual $\|\Theta - T_r(\Theta)\|/\|\Theta\|$ over $\Theta \in$ **912** ($\|\Theta\| = 1$) for alternating projections (AP) versus a BFGS descent. $\|\Theta - T_r(\Theta)$ is the best rank-n approximation of Θ, computed via SVD. Both methods locate the minimum rank at $r = 7$; for $r \geq 7$ AP reaches machine zero while BFGS plateaus at its gradient tolerance.

r	AP residual	BFGS residual
3	7.559×10^{-1}	7.559×10^{-1}
4	6.547×10^{-1}	6.547×10^{-1}
5	5.345×10^{-1}	5.345×10^{-1}
6	3.780×10^{-1}	3.780×10^{-1}
7	2.916×10^{-15}	6.556×10^{-10}
8	2.799×10^{-15}	5.881×10^{-10}
9	2.822×10^{-15}	1.562×10^{-9}
10	2.783×10^{-15}	3.753×10^{-9}

Rank n	Null space dim. d
1–6	0
7–12	1
13–17	3
18–21	6
22–24	10
25–26	15
27	21
28	36

Table 5.2: Null space dimension $d = \dim(\ker A_n)$ for each rank n in the canonical electric frame $\{e_1, \ldots, e_n\}$. The dimension of the space of coefficients for rank n is generally $133n$, and the null space governs the degrees of freedom in the closure search. Values $d > 0$ indicate existence of potential gauge algebra solutions, while $d = 0$ rules out all non-trivial solutions.

Moving on to the actual results, in the following subsections they are all gathered and presented in a systematic way, with the inclusion of **a new exact scan** not presented before, whose need emerged from the limitations of the BFGS stochastic scan.

5.2.1 Results of the BFGS Scan

This section collects, in one place, the full output of the stochastic (BFGS-based) closure scan of `rankn_gauging_analysis.m` over all ranks $1 \leq n \leq 28$ in the canonical electric frame $V_a = e_a$, and gives a reading key for the quantities it reports. The scan proceeds in two stages for each rank: an **exact** linear stage that computes the null space of the $X_{(MNP)} = 0$ (**912**) constraint, and a **stochastic** nonlinear stage that searches that null space for points at which the gauge generators close into a Lie algebra (Section 4.5.1). Note that the computation for the no-go theorem stopped at the exact stage and never reached the stochastic component, which is why it can be presented as a rigorous proof.

A comprehensive reading key for Table 5.3 can be found in Appendix A.2. The **56** Casimir column lists multiplicities of distinct quadratic-Casimir eigenvalues on the **56** (0^m denotes an m-fold zero eigenvalue; $1 \times m$ a non-zero value of multiplicity m). All ranks 7–28 come from the latest unified

Rank	Kill. sig. (p,q,z)	$I_{\mathbb{Q}}$	**56** zeros / nonzero clusters	Closure res.
1–6	— (empty null space, no gauging)			
7	$(0,0,7)$	0	0^{56}	2.2×10^{-17}
8	$(0,0,8)$	0	0^{56}	2.4×10^{-17}
9	$(0,0,9)$	0	0^{56}	1.6×10^{-17}
10	$(0,0,10)$	0	0^{56}	1.8×10^{-17}
11	$(0,0,11)$	0	0^{56}	1.1×10^{-17}
12	$(0,0,12)$	0	0^{56}	1.6×10^{-17}
13 (+)	$(1,0,12)$	$+1/50$	$0^{32},\ 1{\times}24_{\neq0}$	2.0×10^{-16}
13 (−)	$(0,1,12)$	$-1/68$	$0^{32},\ 1{\times}24_{\neq0}$	1.7×10^{-15}
14 (+)	$(1,0,13)$	$+1/49$	$0^{32},\ 1{\times}24_{\neq0}$	1.8×10^{-16}
14 (−)	$(0,1,13)$	$-1/49$	$0^{32},\ 1{\times}24_{\neq0}$	1.1×10^{-15}
15 (+)	$(1,0,14)$	$+1/49$	$0^{32},\ 1{\times}24_{\neq0}$	1.9×10^{-16}
15 (−)	$(0,1,14)$	$-1/52$	$0^{32},\ 1{\times}24_{\neq0}$	1.5×10^{-15}
16 (+)	$(1,0,15)$	$+1/49$	$0^{32},\ 1{\times}24_{\neq0}$	1.8×10^{-16}
16 (−)	$(0,1,15)$	$-1/61$	$0^{32},\ 1{\times}24_{\neq0}$	1.2×10^{-15}
17 (+)	$(1,0,16)$	$+1/49$	$0^{32},\ 1{\times}24_{\neq0}$	1.7×10^{-16}
17 (−)	$(0,1,16)$	$-1/49$	$0^{32},\ 1{\times}24_{\neq0}$	1.2×10^{-15}
18	— (null space non-empty; no closed algebra found)			
19	$(2,1,16)$	$+1/127$	$0^{20},\ 3{\times}12_{\neq0}$	4.5×10^{-4}
20	— (null space non-empty; no closed algebra found)			
21	$(2,1,18)$	$+1/49$	$0^{20},\ 3{\times}12_{\neq0}$	8.7×10^{-4}
22–27	— (null space non-empty; no closed algebra found)			
28 $\mathfrak{so}(5,3)$	$(15,13,0)$	$+1/49$	—	4.4×10^{-4}
28 $\mathfrak{so}(4,4)$	$(16,12,0)$	$+1/50$	—	5.2×10^{-4}
28 $\mathfrak{so}^*(8)\cong\mathfrak{so}(6,2)$	$(12,16,0)$	$+1/178$	—	6.9×10^{-4}

Table 5.3: Results of the BFGS scan from the latest run executed. The run was improved by implementing a multiseed test (seeds used: {12345, 23456, 34567}) in order to retrieve more gaugings and have a more comprehensive scan. Ranks 7–12 each have a one-dimensional null space, yielding a unique solution directly (no BFGS descent): an abelian $\mathbb{R}^n$ gauge algebra with vanishing Killing form and Dynkin index, hence a trivial (all-zero) **56** Casimir spectrum. The name of the Levi factor $\mathfrak{g}_{\rm ss}$ is read off intrinsically from its dimension and Killing signature: ranks 7–17 have a trivial Levi ($\dim\mathfrak{g}_{\rm ss}=0$, abelian or fully solvable), ranks 19 and 21 carry $\mathfrak{g}_{\rm ss}=\mathfrak{sl}(2,\mathbb{R})$ (the split $\dim 3$ form, signature $(2,1)$), and rank 28 is the semisimple $\mathfrak{so}(8)$ real form fixed by (p,q,z). `LieART` supplies the irrep. tables, giving $\mathbf{56}\to 32\,\mathbf{1}\oplus 12\,\mathbf{2}$ for the $\mathfrak{sl}(2,\mathbb{R})$ Levi (ranks 19, 21) and $\mathbf{56}\to\mathbf{28}\oplus\mathbf{28}'$ for the $\mathfrak{so}(8)$ real forms (rank 28).

scan. Despite efforts, BFGS did not prove efficient in the search for gaugings of higher ranks, probably due to the fact that these solutions lie in very narrow basins. Ranks 22–27 reported problems: they are marked "no closed algebra found", despite having a non-trivial null space at the linear stage, and for rank 28 the gaugings SO(8) and SO(1, 7) were never found. This is why the exact scan in Section 4.6 was implemented tentatively, and ultimately it has been able to provide interesting results. A reading key to interpret each quantity in this and the following result tables can be found in Appendix A.2.

5.2.2 What the BFGS scan establishes and what it leaves open

Three qualitative conclusions survive at the level of the stochastic scan. First, the first non-abelian algebra appears at rank 13. The subsequent intermediate ranks $13 \leq n \leq 21$ carry only Levi-decomposable algebras with a large solvable radical ($\dim \mathfrak{r} = 12$ for $n \leq 17$, 18 for $n \geq 18$). In the latest run, every one of them is a genuine gauging, with a non-zero Dynkin index, the $\mathfrak{sl}(2)$-like `Kill1+0` branch sitting at $I \simeq 1/49$. Second, the only semisimple gaugings appear at rank 28, where the scan recovers real forms of $D_4 = \mathfrak{so}(8, \mathbb{C})$, the expected $\mathcal{N} = 8$ SO(p, q)-type gaugings.

From the scan, a principal weakness emerges: due to the stochastic nature of the search, the set of solutions it returns depends on the random seeds, which may or may not correspond to profitable starts. Two symptoms are already visible in the table of Section 5.2.1: Ranks 18, 20 and 22–27, despite a non-trivial null space, solutions are absent for this particular run, and the same result was found for previous, older runs. At rank 28, the dedicated search returns three of the five non-degenerate real forms, $\mathfrak{so}(5, 3)$, $\mathfrak{so}(4, 4)$, $\mathfrak{so}^*(8) \cong \mathfrak{so}(6, 2)$, while missing the compact $\mathfrak{so}(8)$ and the $\mathfrak{so}(7, 1)$ forms entirely. It must be stressed that this scan does not exclude solutions, in the sense that their absence does not imply that they are forbidden. The most probable cause for their absence is that they occupy narrow basins of attraction whose volume is small relative to the sampled region, so a finite number of random starts simply fails to land in them. In other words, the stochastic scan establishes **existence** where it succeeds, but its silence is not a proof of absence. This is exactly why the exact scan was built. We now report the results of this approach in the following section.

5.2.3 Exact scan results: the CSO(p, q, r) family

The procedure reproduces the full CSO(p, q, r) (Appendix A.3) classification documented in [6], but without previous assumption that exploit the particular symmetry of these groups to obtain results, as was done for the BFGS scan, in order to test how it behaves and realize its strengths and weaknesses. For every rank $7 \leq n \leq 28$, closure was certified from the $\delta = 0$ condition explained in Section 4.6. The top (maximal–dimension) families are reported in the following table:

Table 5.4: Results of the exact scan. All families are reported specifying their λ support. Differing from the reference [6], here the CSO(p, q, r) groups are not repeated if they correspond to isomorphic algebras differing simply by orientation (e. g.: (1,2,5) and (2,1,5) in the same family)

ranks	λ (nonzero support coords)	CSO(p, q, r)
7–12	$a\, e_1$	$(1, 0, 7)$
13–17	$a\, e_1 + b\, e_2 + c\, e_3$	$(2, 0, 6), (1, 1, 6)$
18–21	$a\, e_2 + b\, e_4 + c\, e_6$	$(3, 0, 5), (2, 1, 5)$
22–24	$a\, e_2 + b\, e_5 + c\, e_8 + d\, e_{10}$	$(4, 0, 4), (3, 1, 4), (2, 2, 4)$
25–26	$a\, e_2 + b\, e_6 + c\, e_{10} + d\, e_{13} + f\, e_{15}$	$(5, 0, 3), (4, 1, 3), (3, 2, 3)$
27	$a\, e_2 + b\, e_7 + c\, e_{12} + d\, e_{16} + f\, e_{19} + g\, e_{21}$	$(6, 0, 2), (5, 1, 2), (4, 2, 2), (3, 3, 2)$
28	$\sum a_i e_{\{2,9,16,22,27,31,34,36\}}$	$(8, 0, 0), (7, 1, 0), (6, 2, 0), (5, 3, 0), (4, 4, 0)$

with the lower–dimensional sub–flats at each rank realising the contracted descendants (e.g. $(2, 0, 6)$, $(1, 1, 6)$). Within a family the compact form sits at the all–equal–sign endpoint and each sign flip transfers one unit from p to q.

The semisimple, non–contracted members are precisely the rows with $r = 0$: there CSO$(p, q, 0) \cong$

$\mathfrak{so}(p, q)$, so this column reproduces the full chain of real forms of $\mathfrak{so}(8)$. The compact form sits at the all–equal–sign vertex of its λ–family, $\mathrm{CSO}(8, 0, 0) \cong \mathfrak{so}(8)$, and each sign flip of a free parameter λ_p transfers one unit from p to q, reaching the noncompact split forms $\mathfrak{so}(7, 1)$, $\mathfrak{so}(6, 2)$, $\mathfrak{so}(5, 3)$, $\mathfrak{so}(4, 4)$. Furthermore, the scan confirms that the smallest algebra found is at rank 7, and corresponds to the fully abelian $\mathbb{R}^7$ contraction.

5.2.4 Comparison with the BFGS scan

The standard BFGS search returns one local optimum per random start, and as explained coverage of the variety is only statistical, which means that rare components are easily missed. Enumerating the maximal flats S describes entire families at once and, because each component is linear, represents the closure locus exactly rather than by isolated points. This search was able to find the gaugings that were hidden from the BFGS scan, in particular for ranks 18, 20, and 22-27. The two approaches are complementary: BFGS is a quick existence probe at fixed rank, whereas the flat enumeration provides the complete classification used above. Furthermore, the results provided by the BFGS scan correspond to a set of solutions found also by the exact scan, as one can see comparing the invariants computed.

5.3 CSO(p, q, r) in SL(8, ℝ)

In the $\mathrm{SL}(8, \mathbb{R})$ frame, $\mathrm{CSO}(p, q, r)$ is a known electric gauging, built from the symmetric 8×8 matrix [11]:

$$\Theta_{AB}{}^{C}{}_{D} = \delta^{C}_{[A}\,\eta_{B]D}, \qquad \eta = \mathrm{diag}(\underbrace{+1, \ldots}_{p}, \underbrace{-1, \ldots}_{q}, \underbrace{0, \ldots}_{r}), \quad p + q + r = 8\,. \tag{5.6}$$

In fact, this matches perfectly the reduced structure of the embedding tensor in this electric frame. The linear constraint forces the embedding tensor to sit in the $\mathbf{912}$ of $\mathrm{E}_{7(7)}$; decomposing under $\mathrm{SL}(8, \mathbb{R}) \subset E_{7(7)}$, where $\mathbf{56} = \mathbf{28} \oplus \mathbf{28}'$, the $\mathbf{912}$ splits as

$$\mathbf{912} \;\to\; \underbrace{\mathbf{36}}_{\mathrm{Sym}^2\mathbf{8}} \oplus \underbrace{\mathbf{36}'}_{\mathrm{Sym}^2\mathbf{8}^*} \oplus \mathbf{420} \oplus \mathbf{420}'\,. \tag{5.7}$$

The two pieces $\mathbf{36}$ and $\mathbf{36}'$ correspond to symmetric 8×8 matrices η_{AB} and η^{AB} respectively. Requiring the charges to lie in the electric Lagrangian block $V_a = e_a$ (the $\Lambda^2\mathbf{8}$ half, mutually isotropic) kills the $\mathbf{420} + \mathbf{420}'$ and the magnetic $\mathbf{36}'$. What survives is a single $\mathbf{36}$: one symmetric tensor $\eta_{AB} \in \mathrm{Sym}(8)$, exactly (5.6). The associated X_M are block-diagonal in the $\mathbf{28} \oplus \mathbf{28}'$ split, so these models, including $\mathrm{SO}(8)$, $\mathrm{SO}(1, 7)$ and all CSO contractions, are genuinely electric in the $\mathrm{SL}(8, \mathbb{R})$ frame. Basically, an accurate choice of η reconstructs all gaugings found and is equivalent to the appropriate choice of what we called λ for the exact scan. The fact that the first contraction possible $\mathrm{CSO}(1, 0, 7)$ happens exactly at rank 7 is a consistency validation of the theorem for the smallest possible rank. Having the possibility to compare the results of the two methods employed, stochastic and exact, with known literature, made the validity and limitations of these approaches clear.

An interesting note can be made on the null space dimension: its nonzero values follow a pattern of triangular numbers, such that $T_k = \binom{k}{2} = \frac{k(k-1)}{2}$, with $k = 2, \ldots, 7$ for ranks 7–27, and $k = 9$ for rank 28. Fixing the rank to n determines how many eigenvalues of η must be zero: we will call this number r, so $r = 8 - m$ where $m = p + q$ is the size of the non-degenerate block, $m \times m$, which is free. The independent entries of a symmetric $m \times m$ matrix are its $\frac{m(m+1)}{2} = \binom{m+1}{2}$ upper-triangular entries, and this is precisely the null space dimension d.

Concretely, the observed values correspond to $k = m + 1 = 9 - r$ in $T_k = \binom{k}{2}$: The value $k = 8$

r (zero eigenvalues)	$m = 8 - r$	$k = m + 1$	$d = \binom{k}{2}$	base rank $28 - \binom{r}{2}$
0	8	9	36	28
2	6	7	21	27
3	5	6	15	25
4	4	5	10	22
5	3	4	6	18
6	2	3	3	13
7	1	2	1	7

Table 5.5: Triangular-number structure of the null space dimension d as a function of the number r of zero eigenvalues of the symmetric matrix η. The base rank column gives the leading rank at which each free-parameter count first appears.

(which should give $d = 28$) is missing and it corresponds to $r = 1$, and one might ask themselves why are they excluded. Let us first consider the base rank formula, which follows from counting the nonzero generators $X_{[AB]}$ directly. With m nonzero and $r = 8 - m$ zero eigenvalues of η, there are $\binom{m}{2}$ rotation generators, considering both indices in the nonzero block, and $m\,r$ translational generators, where one index is in each block. Considering these two contributions,

$$n \;=\; \binom{m}{2} + m\,r \;=\; \frac{m(m-1)}{2} + m(8-m) \;=\; 28 - \binom{r}{2}, \tag{5.8}$$

which is the base rank entry in Table 5.5. For $r = 1$, since the base-rank formula gives 28, just like for $r = 0$, no independent $d = 28$ family arises for this particular case, that is therefore absorbed.

5.4 Frames and the mass basis

In this section we use the framework developed in [7] and [8] to conduct a brief analysis on group contractions and the mass basis, with the aim of understanding how switching masses on and off might impact the resulting gaugings and rank of the resulting gauge group algebra. In the frame of choice, considering (5.6), the four gravitino masses are the diagonal of η, grouped in four pairs $(\theta_1\theta_1, \ldots, \theta_4\theta_4)$; a vanishing mass annihilates a pair $(r \to r + 2)$ and the signs fix the real form $\mathfrak{so}(p, q)$. However, the source file `Thetaxyz` for this analysis uses a rotated electric frame in $\mathrm{SU}^*(8)$. In this file, the embedding tensor is stored, with entries depending on seven parameters: four mass parameters m_i and three deformation moduli x_i, with $x_i = 1$ the regular reference point corresponding to the $\mathrm{SO}^*(8)$ gauging, and $x_i \to 0$ a singular contraction. By rotating one frame into the other, as will be further explained, the conclusions for this frame will be transferrable to our usual frame.

5.4.1 Numerical verification

All checks use generic masses $m = (\frac{23}{3}, \frac{41}{5}, \frac{59}{7}, \frac{71}{11})$, all coprime numbers to avoid accidental cancellations. Now, `Thetaxyz` is electric in some frame, but not in $\mathrm{SL}(8,\mathbb{R})$. At $x = (1,1,1)$ we computed the column space V of `Thetaxyz` and its restriction of Ω:

quantity	value
$\dim V$	28
$\max\lvert\, \Omega\rvert_V\,\rvert$	3.5×10^{-16}
weight of V in $\mathrm{SL}(8,\mathbb{R})$ electric $\mathbf{28}$	14
weight of V in $\mathrm{SL}(8,\mathbb{R})$ magnetic $\mathbf{28}'$	14
magnetic fraction	0.5

Table 5.6: Column-space V of `Thetaxyz` and its electric/magnetic decomposition in the $\mathrm{SL}(8,\mathbb{R})$ frame, at $x = (1,1,1)$ with generic masses.

`Thetaxyz` is electric specifically in the $\mathrm{SU}^*(8)$ frame. V is split evenly between the electric and magnetic halves of the $\mathrm{SL}(8,\mathbb{R})$ polarisation, so it is not electric in $\mathrm{SL}(8,\mathbb{R})$. Constructing Θ directly from (5.6) in the $\mathrm{SL}(8,\mathbb{R})$ frame yields, for each mass pattern, a gauge generator with identically vanishing off-diagonal block and closure at machine precision.

mass pattern	rank	off-diag block	Killing $\{p,q,z\}$	closure
all nonzero	28	0	$\{0,28,0\}$ ($\mathfrak{so}(8)$)	10^{-16}
one zero	27	0	$\{0,15,12\}$	10^{-16}
two zero	22	0	$\{0,6,16\}$	10^{-16}
three zero	13	0	$\{0,1,12\}$	10^{-16}
all zero	0	0	$\{0,0,0\}$	0

Table 5.7: Rank and Killing signature of $\mathrm{CSO}(p,q,r)$ gaugings constructed directly in the $\mathrm{SL}(8,\mathbb{R})$ frame via (5.6), for equal-sign mass patterns. The rank series 28/27/22/13/0 is frame-independent and matches the dyonic `Thetaxyz` result.

A check was performed to see that the rank column reproduces the same series of the dyonic `Thetaxyz` scan, since rank is frame independent. As already known, therefore, conclusions on rank and other frame independent objects reached in $\mathrm{SU}^*(8)$, such as the algebras and Killing form and signatures, are valid also in our $\mathrm{SL}(8,\mathbb{R})$ frame.

Contraction $x_i \to 0$ of `Thetaxyz`. Following the convention developed in [8], we multiply Θ by the vanishing moduli, which means rescaling the gauge coupling, and taking the limit $x \to 0$ we obtain:

contraction	rank	Killing $\{p,q,z\}$	resulting model
base $x=(1,1,1)$	28	$\{12,16,0\}$	$\mathrm{SO}^*(8)$, dim 28
$x_1 \to 0$	28	$\{0,0,12\}$	$(\mathrm{SO}(4)\times\mathrm{SO}(2,2)) \ltimes T^{16}$
$x_1, x_2 \to 0$	26	$\{0,0,2\}$	$\mathrm{U}(1)^2 \ltimes N^{26}$
$x_1, x_2, x_3 \to 0$	**25**	$\{0,0,1\}$	$\mathrm{CSS}_{N=0} = \mathrm{U}(1) \ltimes T^{24}$

Table 5.8: Result of contractions of the deformation parameters x_i. These results are compatible with those in [8].

This is the contraction $\mathrm{SO}^*(8) \to \cdots \to \mathrm{CSS}$ obtained in [8]. The full limit $x_1 \sim x_2 \sim x_3 \to 0$ lands on $\mathrm{CSS}_{N=0} = \mathrm{U}(1) \ltimes T^{24}$, of dimension $1+24=\mathbf{25}$. The CSS corresponds to a Cremmer–Scherk–Schwarz gauging, fundamental in the study of $\mathcal{N}=8$ supergravity, particularly in the context of spontaneous supersymmetry breaking on flat Minkowski backgrounds and subsequent vacua study. By tracing the whole sequence of contractions, it is shown that the CSS gauging is not just an isolated solution, but the boundary limit of a continuous moduli space of $\mathrm{SO}^*(8)$ and CSO^* gaugings.

Mass-rank analysis of the CSS endpoint. Starting from the rank-25 CSS embedding tensor, we apply the same test for the masses used for $\mathrm{SO}^*(8)$: masses are set to zero one at a time. Patterns of equal masses are also analyzed, while for the non-contracted analysis they proved non-relevant and produced no variations in ranks. The results are in Table 5.9.

zeroed	nonzero mass pattern		CSS rank	$\mathrm{SO}^*(8)$ rank
0	(a,b,c,d)	all distinct	**25**	28
	(a,a,b,c)	2 equal + 2 distinct	23	28
	(a,a,b,b)	2 pairs	21	28
	(a,a,a,b)	3 equal + 1 distinct	19	28
	(a,a,a,a)	all equal	13	28
1	$(a,b,c,0)$	3 distinct	25	27
	$(a,a,b,0)$	2 equal + 1 distinct	23	27
	$(a,a,a,0)$	3 equal	19	27
2	$(a,b,0,0)$	2 distinct	21	22
	$(a,a,0,0)$	2 equal	19	22
3	$(a,0,0,0)$		13	13
4	$(0,0,0,0)$		0	0

Table 5.9: Rank of the CSS and $\mathrm{SO}^*(8)$ embedding tensors as a function of the number of zeroed masses and the degeneracy pattern of the nonzero masses (a,b,c,d all distinct). For $\mathrm{SO}^*(8)$ the rank depends only on the number of zeros; for CSS it decreases monotonically with mass degeneracy within each zero-count stratum.

Two features distinguish the CSS case from $\mathrm{SO}^*(8)$: first of all, zeroing one mass does not lower the rank, which remains stable at 25, unlike $\mathrm{SO}^*(8)$ where a single zero immediately reduces rank by 1. Secondly, equal masses in this case impose an additional degeneracy: all four masses equal gives rank 13, the same as the single-mass case, because the antisymmetric combinations $m_i - m_j$ that contribute independent directions all vanish.

Finally, the rotation matrix $E_{\mathrm{SL}(8,\mathbb{R})\to\mathrm{SU}^*(8)}$ from the $\mathrm{SL}(8,\mathbb{R})$ frame to $\mathrm{SU}^*(8)$ is directly reported from [8] for completion, to show how one would go to make the $\mathrm{SU}^*(8)$ embedding tensor purely electric in our frame:

$$E_{\mathrm{SL}(8,\mathbb{R})\to\mathrm{SU}^*(8)} = \frac{1}{\sqrt{2}}\begin{pmatrix} \mathbb{1}_{12} & & & -\mathbb{1}_{12} & & \\ & \mathbb{1}_{15} & & & \mathbb{1}_{15} & \\ & & 1 & & & 1 \\ \mathbb{1}_{12} & & & \mathbb{1}_{12} & & \\ & -\mathbb{1}_{15} & & & \mathbb{1}_{15} & \\ & & -1 & & & 1 \end{pmatrix}. \tag{5.9}$$

However, a subtle mismatch exists between the $(12, 15, 1)$ block ordering utilized in the literature and our standard lexicographic ordering for the 2-forms $\Lambda^2(\mathbf{8})$. Under our natural basis choice, splitting the fundamental representation as $\mathbf{8} = \{1, 2\} \oplus \{3, \ldots, 8\}$ yields a 2-form decomposition of $1 + 12 + 15$, since the singleton pair $(1, 2)$ appears first lexicographically. To reconcile this, we introduce a permutation matrix P that maps our lexicographic basis into the $(12, 15, 1)$ block structure via:

$$\Lambda^2(\{1,2\}) \oplus (\{1,2\} \otimes \{3,\ldots,8\}) \oplus \Lambda^2(\{3,\ldots,8\}) \xrightarrow{P} (12, 15, 1)\,. \tag{5.10}$$

Applying $E_{\mathrm{lex}}^{-1} = P^\top E^\top P$ we get the following result:

quantity	value	meaning
$\max\lvert E^\top E - \mathbb{1}\rvert$	2.2×10^{-16}	E is orthogonal
$\max\lvert E^\top \Omega E - \Omega\rvert$	2.2×10^{-16}	$E \in \mathrm{Sp}(56, \mathbb{R})$
$\det E$	1.0000	$E \in \mathrm{Sp}(56, \mathbb{R}) \cap \mathrm{O}(56)$
magnetic fraction *before* ($\mathrm{SL}(8, \mathbb{R})$)	0.5000	50% magnetic (dyonic)
magnetic fraction *after* (E_{lex}^{-1} applied)	0 (exact)	purely electric
$\max\lvert(E_{\mathrm{lex}}^{-1}\Theta)_{\mathrm{mag}}\rvert$	0 (exact)	magnetic charges vanish exactly

Table 5.10: Numerical verification of the rotation $E_{\mathrm{lex}}^{-1} = P^\top E^\top P$ applied to `Thetaxyz` at $x = (1, 1, 1)$. The magnetic block vanishes exactly, confirming that `Thetaxyz` is manifestly electric in the $\mathrm{SU}^*(8)$ frame.

The entire magnetic block of $E_{\mathrm{lex}}^{-1}\,\Theta$ vanishes exactly, and one sees the $50\% \to 0$ collapse of the magnetic weight: in the rotated $\mathrm{SU}^*(8)$ frame `Thetaxyz` is manifestly electric.

In this section, we have seen how the masses influence the rank of the gauge group's algebra. By turning them off one by one, the rank gets lower and reaches the lower bound of 13. The mechanism of contractions that lead to a CSS gauging is also fascinating; a full analysis of this and the solid theoretical background for the study of the gauging of flat groups in 4d supergravity is found in the referenced papers [7] and [8].

Chapter 6

Conclusions and outlook

This thesis provides a general overview on the problem of analyzing gaugings of maximal supergravity by employing the rank-stratification approach for the embedding tensor. The main aim was to understand what was the minimum rank resulting in non-trivial gaugings: the multiple verifications presented all agree on the lower bound of $\mathrm{rank}(\Theta) = 7$, establishing a new result that could be used as a basis for further studies. The orbit-invariance of this result has therefore very strong evidence. On the other end, the methods tried to classify gaugings proved limited: both the BFGS scan and the exact scan cannot certify that one has found the whole list of gaugings, however especially the exact scan provides an interesting tool for classification when the scan is characterized by certain conditions (for example, for the in-support case $\delta = 0$). The results are all compatible with the gauging analysis in [6], which served as a check for consistency and validity. The smallest rank-7 gauging is exactly the $\mathrm{CSO}(1, 0, 7)$ solution.

While not claiming to exhaust the matter, therefore, this thesis provides a study on possible ways to approach the rank stratification problem and the minimum rank question, leaving room for further improvements. The next steps in this study would be certainly to improve the reliability of the orbit-independence claim of the minimum support for the linear constraint, since they are subject to computational approximation and sampling. After gauging analysis, another continuation of the study would be the classifications of vacua, particularly employing their dependency on the cosmological constant Λ, but also analyzing their stability properties. Supersymmetry breaking mechanisms provide also an open field in which to investigate. Furthermore, vacua study could then lead to the check for the validity of these theories, characterized by the vacua studied, in the context of String Theory. In particular, the verification of whether these theories can be properly lifted to string theories that are compatible with the Swampland conjectures, is a fascinating question. For example, for the non-SUSY AdS instability conjecture, a consistency requirement for quantum gravity theories is that any Anti-de Sitter vacua that does not possess supersymmetry must be unstable. Ultimately, exploring the algebraic boundaries of maximal supergravity in four dimensions not only uncovers the minimal structural requirements for non-trivial gaugings, but provides us with the right tools to probe which theories are compatible with the Swampland bounds and therefore can be realized as UV-complete string theories.

Appendix A

Appendix

In this section, the notions of linear algebra and group theory needed to fully understand the classification of the gaugings are explained in more detail. The main reference used is [20]. A reading key for Table 5.3 is also provided.

A.1 Definition of some algebraic invariants

Here are reported the definitions for the invariants used in the classification of gaugings. $\mathfrak{g}$ represents the gauge algebra.

1. **Nilpotency**: $\mathfrak{g}$ is nilpotent if the lower central series $\mathfrak{g} \supset [\mathfrak{g}, \mathfrak{g}] \supset [\mathfrak{g}, [\mathfrak{g}, \mathfrak{g}]] \supset \cdots$ reaches $\{0\}$.
2. **Solvability**: $\mathfrak{g}$ is solvable if the derived series $\mathfrak{g}^{(0)} \supset \mathfrak{g}^{(1)} \supset \mathfrak{g}^{(2)} \supset \cdots$ reaches $\{0\}$.
3. **Semisimplicity**: $\mathfrak{g}$ is semisimple if $\operatorname{rank} B = n$.
4. **Unimodularity**: ${f_{ab}}^b = 0$.
5. **E7 Gram matrix**: $G_{ab} = C_a \cdot \texttt{metre7} \cdot C_b^T$. Its signature distinguishes subalgebras with the same abstract type but different embeddings in $\mathfrak{e}_{7(7)}$.

A.2 Reading key

Null dimension d. Output of the exact linear stage: the dimension of $\ker A$. It is 0 for $n \leq 6$ (no gauging is even kinematically possible) and grows with n, from $d = 1$ at $n = 7$ to $d = 36$ at $n = 28$. A non-zero d is necessary but not sufficient for a gauging: the nonlinear closure condition must still be met.

Killing signature (p, q, z). The numbers of positive, negative and zero eigenvalues of the gauge Killing form $K_{ab} = {f_{ac}}^d {f_{bd}}^c$. The zero count z measures how far the algebra is from semisimple: $z = 0$ means semisimple (rank 28), while $z = n - 1$ or $z = n - 3$ signals a large solvable radical (ranks 13–21). The split p versus q encodes the compactness of the semisimple part.

Dynkin index $I_{\mathbb{Q}}$. The rationalised Dynkin index of the **56** under the gauge algebra, $I = \operatorname{tr}_{\mathbf{56}}(C_2)/\operatorname{tr}_{\mathbf{56}}(\sum_{\alpha=1}^{133} t_\alpha^2)$. A non-zero I marks a genuine gauging that acts non-trivially on the scalar/vector sector. Were I to vanish, every eigenvalue of the quadratic Casimir on the

56 would vanish and the gauging would leave the physical spectrum and scalar potential untouched, a phantom gauging. While in previous scans these gaugings appeared due to numerical floating-point errors, in the latest one no phantom solutions of this kind occur: every entry has $I \neq 0$. An exception, of course, are the purely abelian algebras, for which the Dynkin Index is zero but the gaugings are non-phantom.

Branching of the 56 How the fundamental **56** of $\mathrm{E}_{7(7)}$ decomposes under the gauge algebra. `Branching56Evals` is the multiset of quadratic-Casimir eigenvalues on the **56**, reported in the table as multiplicities of its distinct values.

Closure residual. The stochastic stage reports the residual of the bracket against the span of the generators. At ranks 13–17 it is at the machine-precision floor ($\sim 10^{-16}$), so those algebras close essentially exactly. At ranks 19 and 21 it rises to $\sim 10^{-4}$, and similarly at rank 28, reflecting the larger null space and harder optimisation landscape.

A.3 Algebraic Definition of $\mathrm{CSO}(\mathrm{p}, \mathrm{q}, \mathrm{r})$

In the following section, we provide geometric and algebraic description of $\mathrm{CSO}(\mathrm{p}, \mathrm{q}, \mathrm{r})$ groups, and discuss their connection to contractions [6][11]. These groups naturally arise in the context of flux compactifications and gauged supergravities, where the embedding tensor encodes the geometric transitions between different phases of the theory.

The group $CSO(p, q, r)$ (Contracted Special Orthogonal Group) represents a non-semisimple, generalization of the classical pseudo-orthogonal groups $SO(p, q)$.

Let V be a real vector space of dimension $N = p + q + r$. We define a degenerate symmetric metric η_{AB} (with $A, B = 1, \ldots, N$) of the form:

$$\eta_{AB} = \mathrm{diag}(\underbrace{+1, \ldots, +1}_{p}, \underbrace{-1, \ldots, -1}_{q}, \underbrace{0, \ldots, 0}_{r}) \tag{A.1}$$

The Lie algebra $\mathfrak{cso}(p, q, r)$ consists of the generators $M_{AB} = -M_{BA}$ that preserve this degenerate metric structure. The algebraic structure is best understood by splitting the generators according to the blocks of the metric:

- J_{ab}: Generators of the $\mathfrak{so}(p, q)$ subalgebra (where $a, b = 1, \ldots, p + q$).
- T_{ai}: Translation-like generators spanning the coset (where $i = 1, \ldots, r$).

The commutation relations of the algebra are given by [11, 21]:

$$[J_{ab}, J_{cd}] = \eta_{bc} J_{ad} - \eta_{ac} J_{bd} - \eta_{bd} J_{ac} + \eta_{ad} J_{bc} \tag{A.2}$$

$$[J_{ab}, T_{ci}] = \eta_{bc} T_{ai} - \eta_{ac} T_{bi} \tag{A.3}$$

$$[T_{ai}, T_{bj}] = 0\,. \tag{A.4}$$

Structurally, the algebra is therefore non-semisimple, with the semisimple factor $\mathfrak{so}(p, q)$ acting on an abelian translational radical, giving the semidirect product structure

$$\mathfrak{cso}(p,q,r) \cong \mathfrak{so}(p,q) \ltimes \mathbb{R}^{r\cdot(p+q)}. \tag{A.5}$$

The $CSO(p,q,r)$ groups can be formally obtained from smooth, semisimple pseudo-spherical groups such as $SO(p+r,q)$ or $SO(p,q+r)$ through a singular limiting procedure known as an **Inönü–Wigner contraction** [21].

Physically, this process corresponds to flat limit of the underlying geometry, or a scaling limit where certain coupling constants or fluxes are sent to zero.

Consider the original semisimple algebra $\mathfrak{so}(p+r,q)$ spanned by the generators M_{AB}. We split the index A into "solid" indices $a = 1,\ldots,p+q$ and "contracted" indices $i = 1,\ldots,r$. The generators are accordingly reorganized into J_{ab} and K_{ai}.

By defining a parameter-dependent, singular change of basis governed by a real parameter $\epsilon > 0$:

$$J'_{ab} = J_{ab}\,, \qquad T_{ai} = \epsilon\, K_{ai} \tag{A.6}$$

In the parent algebra the mixed generators close on $[K_{ai}, K_{bj}] = -\eta_{ij} J_{ab} - \eta_{ab} M_{ij}$, where M_{ij} generate rotations among the contracted directions. Rescaling and taking the contraction limit $\epsilon \to 0$, both terms are suppressed by ϵ^2:

$$[T_{ai}, T_{bj}] = \lim_{\epsilon\to 0} \epsilon^2 [K_{ai}, K_{bj}] = -\lim_{\epsilon\to 0} \epsilon^2 (\eta_{ij} J_{ab} + \eta_{ab} M_{ij}) = 0\,. \tag{A.7}$$

As a result of this limit, the generators T_{ai} become purely **abelian translations**.

A.4 SU*(2n)

The non-compact real form SU*(2n) is the “quaternionic” real form of A_{2n-1} [22]. For $n = 4$, it appears as one of the electric duality frames of $\mathrm{E}_{7(7)}$, on the same footing as $\mathrm{SL}(8,\mathbb{R})$ [11].

Let J be a real $2n \times 2n$ matrix with $J^2 = -\mathbb{1}_{2n}$, defining a quaternionic structure $j(v) = J\bar{v}$ on $\mathbb{C}^{2n}$ (an antilinear map with $j^2 = -1$). The algebra consists of the traceless complex matrices that commute with j,

$$\mathfrak{su}^*(2n) = \{\, X \in \mathfrak{sl}(2n,\mathbb{C}) \;:\; X = J\,\bar{X}\,J^{-1}\,\}, \tag{A.8}$$

that is, the $n \times n$ quaternionic matrices of vanishing real trace. Being a real form of $\mathfrak{sl}(2n,\mathbb{C})$, it has $\dim_{\mathbb{R}} \mathfrak{su}^*(2n) = 4n^2 - 1$, and for $n = 4$ this gives $\dim \mathfrak{su}^*(8) = 63$.

The SU*(8) **frame and** $CSO^*(2p,2q)$ **gaugings.** $SU^*(8)$ provides an electric frame of $\mathrm{E}_{7(7)}$ such that the resulting branchings under $\mathrm{E}_{7(7)}$ representations take formally the same form as in the $SL(8,\mathbb{R})$ frame [11]. The gaugings built only from the electric vectors $A_\mu^{a'b'}$ are encoded by an embedding-tensor component $\theta_{a'b'}$ in the **36** of $SU^*(8)$, which a suitable $SU^*(8)$ transformation brings to the diagonal form:

$$\theta_{a'b'} = \mathrm{diag}(\underbrace{1,\ldots,1}_{2p}, \underbrace{0,\ldots,0}_{2q}), \qquad p + q = 4. \tag{A.9}$$

The corresponding gauge group CSO*(2p, 2q) is the subgroup of $SU^*(8)$ that preserves $\theta_{a'b'}$. As in the $SL(8,\mathbb{R})$ case, it is non-semisimple and of contracted type,

$$\mathfrak{cso}^*(2p, 2q) \cong \mathfrak{so}^*(2p) \ltimes \mathbb{R}^{2p\cdot 2q}, \tag{A.10}$$

with a semisimple factor $\mathfrak{so}^*(2p)$ acting on an abelian radical of commuting nilpotent translations, constituting the SU* analogue of the CSO(p, q, r) family of Appendix A.3. The $2q$ vanishing entries of θ playing the role of the degenerate block. The semisimple endpoint $q = 0$ gives $CSO^*(8,0) = SO^*(8) \cong SO(6,2)$.

A.5 SO*(2n)

The non-compact real form SO*(2n) is the real form of $D_n = \mathfrak{so}(2n,\mathbb{C})$ that preserves a quaternionic structure. It is the semisimple factor of the $CSO^*(2p,2q)$ gaugings of Appendix A.4, and for $n = 4$ it is one of the rank-28 real forms recovered by the scan.

With the quaternionic structure [22] J ($J^2 = -\mathbb{1}_{2n}$, $j(v) = J\bar{v}$, $j^2 = -1$),

$$\mathfrak{so}^*(2n) = \{\, X \in \mathfrak{so}(2n,\mathbb{C}) \;:\; X = J\,\bar{X}\,J^{-1} \,\}, \tag{A.11}$$

Equivalently, in a block basis it is realised by

$$X = \begin{pmatrix} A & B \\ -\bar{B} & \bar{A} \end{pmatrix}, \qquad A^\dagger = -A, \quad B^T = -B, \tag{A.12}$$

with A anti-Hermitian and B complex antisymmetric. This gives $\dim_{\mathbb{R}} \mathfrak{so}^*(2n) = n(2n-1)$, with maximal compact subalgebra $\mathfrak{u}(n)$ of dimension n^2, so that the quotient $SO^*(2n)/U(n)$ is a Hermitian symmetric space.

The relevant accidental isomorphisms are

$$\mathfrak{so}^*(4) \cong \mathfrak{su}(2) \oplus \mathfrak{sl}(2,\mathbb{R}), \qquad \mathfrak{so}^*(6) \cong \mathfrak{su}(3,1), \qquad \mathfrak{so}^*(8) \cong \mathfrak{so}(6,2). \tag{A.13}$$

For $n = 4$ one has $\dim \mathfrak{so}^*(8) = 28$, and the D_4 isomorphism $\mathfrak{so}^*(8) \cong \mathfrak{so}(6,2)$ identifies it with one of the non-compact real forms of D_4 appearing among the rank-28 gaugings (its split companions being $\mathfrak{so}(5,3)$ and $\mathfrak{so}(4,4)$). It is the semisimple part surviving in CSO*(8, 0).

Bibliography

[1] Michael B. Green, John H. Schwarz, and Edward Witten. *Superstring Theory. Volume 1: Introduction*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 1987.

[2] Eugene Cremmer and Bernard Julia. The $N = 8$ Supergravity Theory. 1. The Lagrangian. *Phys. Lett. B*, 80:48, 1978.

[3] Steven Weinberg and Edward Witten. Limits on Massless Particles. *Phys. Lett. B*, 96:59–62, 1980.

[4] Bernard de Wit, Henning Samtleben, and Mario Trigiante. The maximal D=4 supergravities. *JHEP*, 06:049, 2007.

[5] Gianguido Dall'Agata and Marco Zagermann. *Supergravity*, volume 991 of *Lecture Notes in Physics*. Springer, 2021.

[6] F. Cordaro, P. Fré, L. Gualtieri, P. Termonia, and M. Trigiante. $\mathcal{N} = 8$ gaugings revisited: An exhaustive classification. *Nucl. Phys. B*, 532:245–279, 1998.

[7] Laura Andrianopoli, Riccardo D'Auria, Sergio Ferrara, and María A. Lledó. Gauging of Flat Groups in Four Dimensional Supergravity. *JHEP*, 07:010, 2002.

[8] Francesco Catino, Gianguido Dall'Agata, Gianluca Inverso, and Fabio Zwirner. On the moduli space of spontaneously broken N = 8 supergravity. *JHEP*, 09:040, 2013.

[9] Henning Samtleben. Lectures on Gauged Supergravity and Flux Compactifications. *Class. Quant. Grav.*, 25:214002, 2008.

[10] Gianguido Dall'Agata, Gianluca Inverso, and Alessio Marrani. Symplectic Deformations of Gauged Maximal Supergravity. *JHEP*, 07:133, 2014.

[11] Mario Trigiante. Gauged supergravities. *Phys. Rept.*, 680:1–175, 2017.

[12] E. Cremmer and B. Julia. The so(8) supergravity. *Nucl. Phys. B*, 159:141–212, 1979.

[13] B. de Wit and H. Nicolai. N=8 supergravity. *Nucl. Phys. B*, 208:323–364, 1982.

[14] Bernard de Wit, Henning Samtleben, and Mario Trigiante. On Lagrangians and gaugings of maximal supergravities. *Nucl. Phys. B*, 655:93–126, 2003.

[15] G. Dall'Agata, G. Inverso, and M. Trigiante. Evidence for a family of SO(8) gauged supergravity theories. *Phys.Rev.Lett.*, 109:201301, 2012.

[16] Gianguido Dall'Agata, Gianluca Inverso, and Alessio Marrani. Symplectic Deformations of Gauged Maximal Supergravity. *JHEP*, 07:133, 2014.

[17] Robert Feger, Thomas W. Kephart, and Robert J. Saskowski. LieART 2.0 – A Mathematica Application for Lie Algebras and Representation Theory. *Comput. Phys. Commun.*, 257:107490, 2020.

[18] J. C. A. Barata and M. S. Hussein. The Moore–Penrose pseudoinverse: A tutorial review of the theory. *Brazilian Journal of Physics*, 42(1):146–165, 2012.

[19] G. Dall'Agata and G. Inverso. On the vacua of gauged supergravity in 4 dimensions. *Nuclear Physics B*, 859(1):70–95, 2012.

[20] Pavel Etingof. Lie groups and lie algebras, 2022.

[21] E. Inonu and E. P. Wigner. On the contraction of groups and their representations. *Proceedings of the National Academy of Sciences of the United States of America*, 39(6):510–524, 1953.

[22] Anthony W. Knapp. *Lie Groups Beyond an Introduction.* Published by the Author, East Setauket, New York, digital second edition edition, 2023.