Beyond First-Order Interface Coupling: Second-Order Closure Corrections in the VERSF Framework

Keith Taylor VERSF Theoretical Physics Program

For the General Reader

In physics, some numbers appear fundamental. One of the most famous is the fine-structure constant, which sets the strength of the electromagnetic force. Its inverse is approximately 137 — a number that has puzzled physicists for decades.

In earlier work within the VERSF framework, we showed that this number can be approximated from first principles by analysing the structure of the boundary where physical facts form. That calculation produced a value extremely close to 137 — within about 0.08%.

This paper takes the next step. Rather than treating that small difference as an error, we show that it arises from a deeper layer of structure that the first calculation intentionally simplified. By incorporating interactions between constraints at this boundary — what we call second-order closure correlations — we derive a correction term that brings the result into agreement at the 10^{-3} level with the observed value.

The key insight is that the constraints governing physical interactions at the interface are not independent. They share a single global consistency condition, and this shared condition introduces correlations that slightly reduce the effective strength of the coupling. The interface has 6 local constraints and 1 global closure constraint. Because the global constraint is shared equally among the 6 local ones — a symmetry of the hexagonal interface structure established in prior VERSF work — each local constraint carries the same participation share, and the magnitude of the second-order reduction is determined by this equal sharing rather than by any fitted parameter.

The result is not just a better number. It is evidence that the structure uncovered in the earlier work is not coincidental: it continues to generate accurate physical predictions when examined at the next level of approximation. We also identify explicitly where further derivation is needed to make this a fully closed argument.

Table of Contents

---

1. Abstract

The VERSF interface realization programme previously derived a leading-order structural expression for the electromagnetic coupling,

$g\_eff^{-2} = 2^K \cdot (N+1)/N$,

which evaluates to 137.143 for K = 7, N = 14, matching the observed inverse fine-structure constant to within 0.08%.

In this paper we show that this residual discrepancy arises from a mean-field approximation embedded in the first-order treatment: the global closure constraint is modelled as uniformly redistributed across N independent channels, whereas in reality it couples those channels and suppresses their joint contributions. Relaxing this assumption and incorporating the second-order correlations induced by the shared closure condition yields a structurally determined correction term.

The interface consists of K = 7 constraints: 6 local constraints and 1 global closure constraint. Under the assumption that the global constraint is distributed uniformly across the 6 local degrees of freedom — a symmetry to be derived from the VERSF hexagonal interface architecture in future work — the inverse participation ratio (IPR) equals 1/6. This quantity, combined with a constraint competition model that gives the correction a negative sign and normalisation Var(G) = 1 for the global mode, yields the corrected expression:

$g^{-2} = 2^K \cdot [(N+1)/N - 1/(6N^2)]$

For $K = 7$, $N = 14$, this gives $g^{-2} \approx 137.034$, in agreement with the observed value $\alpha^{-1} \approx 137.036$ at the 10^{-3} level, significantly improved relative to first order. The result is conditional on four structural assumptions stated explicitly in Section 5.3. The identification of g^{-2} with the fully renormalized infrared coupling $\alpha^{-1}(0)$ also remains to be derived within the framework, and the numerical agreement may partly reflect cancellation between this unresolved identification and higher-order corrections.

---

2. Introduction

The VERSF reconstruction programme has established that the structural requirements of stable fact formation constrain the architecture of physical law. The gauge-invariant interface at which physical commitments form is characterised by a constraint structure that, at leading order, yields a predictive expression for the electromagnetic coupling:

$g\_eff^{-2} = 2^K \cdot (N+1)/N$.

With $K = 7$ and $N = 14$ this gives:

$g\_eff^{-2} = 128 \cdot (15/14) \approx 137.143$.

The observed inverse electromagnetic coupling at low energy is $\alpha^{-1} \approx 137.036$, leaving a residual discrepancy of:

$\Delta \approx 0.107$ (~0.08%).

This paper investigates the structural origin of this residual. We argue that it arises from a well-defined approximation made at first order — the assumption of channel independence under the global closure constraint — and derive the leading correction from the constraint structure of the interface. We state explicitly which steps are fully derived and which rest on structural assumptions not yet derived within the framework.

---

3. First-Order Approximation and Its Limitation

The redistribution factor $(N+1)/N$ emerges from modelling the global closure constraint as uniformly redistributed across N independent channels. This is a mean-field approximation: each channel is treated as contributing independently to the overall closure probability, with the global constraint entering only as a uniform shift.

This approximation is consistent at leading order but has a structural limitation. The global closure constraint is not independently applied to each channel; it is a single shared condition that all channels must collectively satisfy. When one channel's local constraint is satisfied, it draws on a shared capacity, reducing what is available to others. The channels are therefore coupled through the global closure condition as a shared resource, and this coupling is invisible at first order.

The first-order result is therefore expected to be accurate only up to corrections of order $1/N^2$, arising from pairwise correlations between channels induced by the shared constraint. The observed residual $\Delta \approx 0.107$ is consistent with this expectation:

$$1/N^2 = 1/196 \approx 5.1 \times 10^{-3}, \quad 128 \cdot (1/N^2) \approx 0.653.$$

The correction is of order $1/N^2$ but requires a structural coefficient and sign to determine its precise form.

---

4. Constraint Structure of the Interface

The interface is characterised by $K = 7$ total constraints, decomposed as:

- **6 local constraints**, each governing an independent local degree of freedom;
- **1 global closure constraint**, enforcing collective consistency of the entire interface.

Uniformity assumption. We assume that the global constraint is distributed equally across all 6 local degrees of freedom, with participation weights:

$$w_i = 1/6 \quad \text{for } i = 1, \ldots, 6.$$

This uniform distribution is a structural symmetry of the VERSF hexagonal interface. In the hexagonal interface architecture, the 6 local constraints correspond to the 6 edges of the fundamental hexagonal cell, and the global closure constraint enforces overall planarity. The 6-fold rotational symmetry of this structure — established in prior VERSF work on the hexagonal interface geometry — implies that no local constraint is distinguished from any other with respect to the global closure condition, forcing equal participation weights. A full derivation of this symmetry from the VERSF field equations is deferred to a companion paper; for the present purposes it is stated as a structural assumption.

Each local constraint therefore carries an equal participation share of 1/6 in the one-dimensional global closure subspace. The remaining global-subspace participation is shared across the other local constraints through the same shared subspace, thereby inducing non-independence between them.

---

5. Derivation of the Second-Order Correction

5.1 — Constraint Participation

Each closure event requires the simultaneous satisfaction of 6 local constraints and 1 global closure constraint. The global constraint contributes one shared degree of freedom distributed equally across all 6 local constraints under the uniformity assumption of Section 4. Without this coupling, each local constraint would have an independent satisfaction probability. With it, the effective participation share of each local constraint in the global closure mode is 1/6 — the reciprocal of the number of local constraints sharing the single global condition.

5.2 — Structure of Pairwise Correlations

The closure probability may be expanded as:

$P\_closure = P^{(1)} + P^{(2)} + \cdots$

where $P^{(1)}$ encodes independent single-channel contributions and $P^{(2)}$ encodes leading-order pairwise corrections.

At first order, each of the N channels contributes a term of order 1/N, and their sum yields the redistribution factor (N+1)/N. At second order, pairwise contributions between channels i and j scale as $1/N^2$, reflecting the joint probability of two-channel contributions.

A critical distinction must be made. While all N channels depend on the global closure constraint, universal dependence is not the same as pairwise correlated suppression. The second-order correction arises specifically from the reduction in joint closure capacity when two local constraints are simultaneously satisfied. This reduction is mediated by their shared dependence on the global constraint as a finite capacity resource. The second-order term is not a sum over all C(N,2) pairs with equal weight — it is governed by the pairwise competition induced by the shared global constraint, whose coefficient is the IPR of the global constraint's distribution, established in Lemma 5.3 below.

5.3 — Lemma 5.3: IPR and the Second-Order Correction

The derivation rests on four structural assumptions, stated here before the proof.

Assumption A (Uniformity). The global constraint is distributed uniformly across the 6 local constraints with weights $w_i = 1/6$. This is the symmetry assumption of Section 4, not yet fully derived from first principles within VERSF.

Assumption B (Constraint competition). The global closure constraint acts as a shared capacity: satisfying local constraint i uses fraction w_i of the global capacity, suppressing joint satisfaction of i and j relative to independent satisfaction, and giving a negative pairwise correlation. This competition model is the physical content of the sign of the correction.

Assumption C (Normalisation). The global constraint fluctuation G is normalised so that $\mathrm{Var}(G) = 1$. This is a convention consistent with treating G as a binary (satisfied/unsatisfied) constraint at the interface level.

Assumption D (Per-channel self-competition). The second-order correction to R is defined as the sum of per-channel self-competition variances — the variance each channel k accumulates through its own coupling w_k to G, computed independently of the other channels — rather than the full variance of the aggregate rate R_G. This choice reflects the interpretation of the correction as the residual loss of independent closure capacity attributable to each local constraint's own occupation of the shared global mode, rather than the total fluctuation of the averaged rate. The physical motivation is developed in Step 3 below.

Lemma 5.3. *Under Assumptions A, B, C, and D, the global closure constraint distributed uniformly across 6 local constraint degrees of freedom has inverse participation ratio IPR = 1/6, and the resulting second-order correction to the redistribution factor R is* $\delta R = -1/(6N^2)$*. This correction is the unique per-channel self-competition term at second order.*

Proof.

Step 1 — Inverse participation ratio. For a normalised weight distribution $\{w_1,\ldots,w_n\}$ with $\Sigma_i w_i = 1$, the IPR is:

$\mathrm{IPR} = \Sigma_i w_i^2$

Under Assumption A, $w_i = 1/6$ for each $i = 1,\ldots,6$, giving:

$\mathrm{IPR} = \Sigma_i (1/6)^2 = 6 \times (1/36) = 1/6$. [sum over i = 1 to 6]

The IPR is uniquely determined by the uniform weight distribution and requires no normalisation choice beyond Assumption A.

Step 2 — Pairwise covariance from constraint competition. Under Assumptions B and C, each local variable X_k has a G-induced component $X_{k,G} = w_k \cdot G$ with $w_k = 1/6$. Because G is a capacity constraint, the pairwise covariance between the G-induced components of channels i and j is:

$\mathrm{Cov}(X_{i\_}G, X_{j\_}G) = -w_i \cdot w_j \cdot \mathrm{Var}(G) = -1/36$.

The negative sign follows from Assumption B; the magnitude from the coupling weights and Assumption C.

Step 3 — Per-channel self-competition and the decomposition. The redistribution factor $R = (1/N)\, \Sigma_k X_k$ is a per-channel average. Its full variance from G is:

$\mathrm{Var}(R\_G) = (1/N^2)(\Sigma_k w_k)^2\, \mathrm{Var}(G) = 1/N^2$.

This decomposes via the algebraic identity $(\Sigma_k w_k)^2 = \Sigma_k w_k^2 + \Sigma_{i \neq j} w_i w_j$ as:

$\mathrm{Var}(R\_G) = \mathrm{IPR}/N^2 + (1 - \mathrm{IPR})/N^2$.

The first term is the sum of per-channel self-competition variances:

$(1/N^2)\, \Sigma_k \mathrm{Var}(w_k \cdot G) = (1/N^2)\, \Sigma_k w_k^2 \cdot \mathrm{Var}(G) = \mathrm{IPR}/N^2$.

The second term is the sum of cross-channel covariance terms:

$(1/N^2)\, \Sigma_{i \neq j} \mathrm{Cov}(w_i \cdot G, w_j \cdot G) = (1 - \mathrm{IPR})/N^2 = 5/(6N^2)$.

Note that $(1 - \mathrm{IPR})/N^2$ is the cross-term remainder in the square of the first-order participation sum $(\Sigma_k w_k)^2$. On physical grounds (Assumption D), this cross-channel piece is excluded from the second-order correction: the cross-channel correlations it encodes are already the content of the mean-field redistribution at first order, and retaining them at second order would reintroduce information already captured. The second-order correction therefore retains only the per-channel self-competition term, with the competition sign from Assumption B:

$\delta R = -\mathrm{IPR}/N^2 = -1/(6N^2)$.

Step 4 — Uniqueness under Assumption D. The per-channel self-competition sum involves $\Sigma_k w_k^2 = \mathrm{IPR}$ at second order. The first-order mean uses $\Sigma_k w_k = 1$ (already absorbed). The third-order self-competition involves $\Sigma_k w_k^3 = 1/36$, entering at $1/N^3$. Under Assumption D, $\delta R = -\mathrm{IPR}/N^2$ is the unique second-order per-channel self-competition correction. ■

Remark. The result is conditional on Assumptions A–D. An earlier version of this paper attempted to derive Assumption D from cumulant theory, but that argument contained an algebraic error: the connected second cumulant of R_G as a whole is $\mathrm{Var}(R\_G) = 1/N^2$, not IPR/N^2, so standard cumulant theory applied to R does not select the diagonal term. Assumption D is therefore stated as an independent physical assumption: the second-order correction is modelled as per-channel self-competition, with cross-channel correlations belonging to the first-order mean-field. This is physically motivated but not yet derived from VERSF first principles.

5.4 — Corrected Redistribution Factor

By Lemma 5.3, the leading second-order correction to R is:

$\delta R = -1/(6N^2)$.

The corrected redistribution factor is therefore:

$R = (N+1)/N + \delta R = (N+1)/N - 1/(6N^2)$.

5.5 — Corrected Coupling Expression

The corrected coupling expression follows directly:

$$\| \; g^{-2} = 2^{K} \cdot [\, (N+1)/N - 1/(6N^{2}) \,] \; \|$$

Each element has a distinct structural origin:

Term	Origin
2^{K}	Simultaneous satisfaction cost of the K binary interface constraints
$(N+1)/N$	Mean-field closure redistribution (first-order)
$1/(6N^{2})$	IPR × per-channel self-competition (second-order, Lemma 5.3, Assumptions A–D)

---

6. Numerical Evaluation

Substituting $K = 7$, $N = 14$:

$1/(6N^{2}) = 1/(6 \cdot 196) = 1/1176 \approx 8.503 \times 10^{-4}$

$(N+1)/N = 15/14 \approx 1.071429$

$R \approx 1.071429 - 0.000850 = 1.070578$

$g^{-2} \approx 128 \times 1.070578 \approx 137.034$

The observed value is $\alpha^{-1} \approx 137.036$. The agreement is at the 10^{-3} level, an improvement of approximately 40-fold relative to first order.

Quantity	Value
First-order result	137.143
Second-order result	137.034
Observed α^{-1}	137.036
First-order discrepancy	~0.08%
Second-order discrepancy	~0.002%

---

7. Structural Interpretation

The correction term has a clear structural origin, conditional on the assumptions of Lemma 5.3.

The **first-order term** $(N+1)/N$ reflects the mean-field contribution: N channels contribute independently, and the global closure constraint shifts the effective channel count by one.

The **second-order term** $1/(6N^2)$ reflects per-channel self-competition through the shared global closure mode. Its coefficient 1/6 is the IPR of the global constraint's uniform distribution across the 6 local degrees of freedom — the unique intensive scalar appearing at second order in the per-channel self-competition expansion, fixed entirely by the interface constraint count under the uniformity assumption.

The key structural claim of the paper is this: **the coefficient 1/6 is not fitted to the observed value.** It follows from the IPR of the assumed uniform distribution over 6 local constraints. Whether this constitutes a fully parameter-free prediction depends on whether Assumptions A–D can be derived from first principles within VERSF, which is identified as the primary open problem.

---

8. Status of the Result and Open Problems

The corrected expression is a second-order structural approximation resting on four structural assumptions and subject to two further identification problems. All are stated here.

Assumption A — Uniformity of constraint distribution. The derivation assumes that the global closure constraint is distributed equally across all 6 local constraints ($w_i = 1/6$). This follows from the 6-fold rotational symmetry of the hexagonal interface geometry. Deriving this symmetry rigorously from the VERSF field equations is an open problem and the primary gap in the current argument.

Assumption B — Competition sign. The negative sign of δR follows from treating the global constraint as a shared capacity that local constraints compete for, rather than a shared enabler. This competition model is physically motivated by the VERSF architecture but its derivation from the formal constraint structure has not been completed. If the competition model is incorrect and the sign were positive, the second-order correction would worsen agreement with observation, which would itself be informative about the framework.

Assumption C — Normalisation Var(G) = 1. The amplitude of the global constraint fluctuation is set to 1 by convention, consistent with a binary constraint at the interface level. This convention is consistent with prior VERSF work but is not independently derived here.

Assumption D — Per-channel self-competition. The second-order correction is modelled as the sum of per-channel self-competition variances IPR/N^2, rather than the full variance of R which gives $1/N^2$. This choice reflects the interpretation of the correction as the residual loss of independent closure capacity attributable to each local constraint's own occupation of the shared global mode, rather than the total fluctuation of the averaged rate. The decomposition identity $(\Sigma_k w_k)^2 = \mathrm{IPR} + (1 - \mathrm{IPR})$ shows that the excluded cross-channel piece $(1 - \mathrm{IPR})/N^2 = 5/(6N^2)$ is five times larger than the retained diagonal piece $\mathrm{IPR}/N^2 = 1/(6N^2)$. The argument for excluding the larger piece is that the cross-term remainder $(1 - \mathrm{IPR})/N^2$ encodes collective cross-channel fluctuations of the same shared mode whose total participation already fixes the mean-field redistribution at first order; under Assumption D these do not define a new per-channel residual and are therefore excluded from the second-order correction. This is the physical motivation for Assumption D; it has not yet been derived from VERSF first principles. An earlier version of this paper claimed this exclusion followed from cumulant theory; that argument contained an algebraic error and has been corrected.

Higher-order corrections. The expansion captures second-order correlations but omits third-order and higher terms. These enter at order $1/N^3$ and are expected to contribute at the level of $\sim 10^{-5}$.

Identification with the infrared coupling. The coupling g^{-2} derived here is a structural quantity defined at the interface level. Its precise identification with the fully renormalized infrared coupling $\alpha^{-1}(0)$ — the quantity measured in low-energy experiments — is not established within the framework. Running coupling effects in standard quantum field theory shift α^{-1} from its ultraviolet to its infrared value by a calculable amount. An analogous derivation within VERSF remains an open problem, and the close numerical agreement may partly reflect cancellation between this unresolved identification and higher-order structural corrections.

The paper's primary claim, accurately stated, is this: conditional on Assumptions A–D and the constraint architecture ($K = 7$, $N = 14$), the second-order correction to the per-channel closure rate is uniquely determined to be $-1/(6N^2)$. This produces numerical agreement at the 10^{-3} level with the observed α^{-1}, an improvement of 40-fold over first order, without fitting any parameter to the observed value.

---

9. Conclusion

The 0.08% discrepancy between the first-order VERSF coupling prediction and the observed inverse fine-structure constant is not a failure of the framework. It is the signature of second-order constraint competition within the interface architecture.

The leading-order model treats N channels as independent under the global closure constraint. This is a mean-field approximation, corrected at second order by per-channel self-competition through the shared global closure mode. Under the uniform distribution assumption (Assumption A) and the competition sign convention (Assumption B), the second-order correction to the per-channel closure rate is $-\mathrm{IPR}/N^2 = -1/(6N^2)$, where $\mathrm{IPR} = 1/6$ is the unique intensive scalar from the weight distribution appearing at this order. The cross-term remainder $(1 - \mathrm{IPR})/N^2$ is excluded under Assumption D on physical grounds: the cross-channel correlations it encodes are already the content of the first-order mean-field redistribution and carry no genuinely new second-order information.

Incorporating this correction yields:

$g^{-2} = 2^K \cdot [\ (N+1)/N - 1/(6N^2)\ ] \approx 137.034,$

in agreement with $\alpha^{-1} \approx 137.036$ at the 10^{-3} level. This agreement is conditional on Assumptions A–D identified in Section 8. The primary open problems — deriving the uniformity symmetry and the competition sign from VERSF first principles, completing the per-channel self-competition model, and identifying g^{-2} with $\alpha^{-1}(0)$ — are clearly identified and constitute the agenda for subsequent work.

The result demonstrates that the VERSF interface framework produces improved physical predictions beyond leading order through a structurally motivated mechanism, and that the constraint architecture encodes genuine second-order physics whose coefficient is determined by the interface geometry rather than by fitting to observation.