

Phase offset ambiguities in multi-echo field mapping

Andrew Van

July 28, 2026

Contents

1	Introduction	1
2	Phase offset estimation	1
2.1	Signal model	2
2.2	Echo spacing	2
2.3	Recovering the offset	2
3	The candidate ladder	3
3.1	The rule ROMEO applies	3
4	Why the rule fails on some frames	4
5	Detecting a failed choice	5
5.1	The signature	5
5.2	What the intercept can and cannot settle	5
6	The selection algorithm	6
7	Estimating the field	7
8	The algorithm in full	7
8.1	What the components must provide	8
8.2	Reference data	8
A	Proof of Proposition 1	9
A.1	The geometry	9
A.2	Consequences	10

1 Introduction

Multi-echo field mapping estimates the field from the rate at which phase accumulates across echoes. Every voxel also carries a receive-coil phase offset, present at zero echo time, common to all echoes, and set by coil geometry rather than by the field. This offset varies rapidly in space, steeply near coil elements and discontinuously between receive channels, so measured phase is not spatially smooth even where the field is. Spatial unwrapping assumes that neighbouring voxels differ by less than half a turn, and raw phase violates that assumption. Removing the offset restores the smoothness later processing needs.

The offset is never measured. It must be inferred from the same wrapped phase that it corrupts, and the inference leaves one integer undetermined. This note sets out the consequences of that integer, the rule conventionally used to fix it, the conditions under which the rule fails, and the procedure MEDIC uses to identify and repair those cases.

2 Phase offset estimation

Notation. A quantity is **wrapped** when confined to $(-\pi, \pi]$.

θ, f		coil offset (rad) and field offset (Hz)
φ_e	wrapped	the recorded phase at echo e
$\Delta\varphi$	wrapped	the wrapped phase difference $W(\varphi_1 - \varphi_0)$
$\Delta\varphi_{\text{uw}}$	unwrapped	that difference with its turn count restored
N	integer	index of a candidate; $N = 0$ is what the unwrapper returned
N^*	integer	index of the correct candidate; not observable
$\hat{\theta}$	wrapped	the offset estimate
ψ_e	neither	recorded phase less an estimated offset

Units. Echo times are in seconds throughout, so a phase slope divided by 2π is a frequency in Hz.

2.1 Signal model

Acquire E echoes at times $t_0 < t_1 < \dots < t_{E-1}$. Field evolution is linear in time, so the phase at echo e before wrapping is $\theta + 2\pi f t_e$. What is recorded is that value wrapped into one turn. With $W(x) = \arg(e^{ix})$,

$$\varphi_e = W(\theta + 2\pi f t_e). \quad (1)$$

Without W , two echoes would determine θ and f by a straight-line fit. With it the line is known only modulo 2π , and infinitely many (θ, f) pairs remain compatible with the data.

Two properties of W are used below. It is invariant under whole turns, $W(x + 2\pi m) = W(x)$ for integer m , so every real x decomposes uniquely as $x = W(x) + 2\pi m$. Wrapping discards m and unwrapping recovers it. Second, a difference of wrapped values returns the wrapped difference only after a second wrap,

$$W(W(x) - W(y)) = W(x - y), \quad (2)$$

because $W(x) - W(y)$ spans $(-2\pi, 2\pi)$ and coincides with $W(x - y)$ only when it happens to land in $(-\pi, \pi]$ already.

2.2 Echo spacing

An echo-planar train produces echoes at constant spacing,

$$t_e = \text{TE}_0 + e \cdot \Delta \text{TE}, \quad e = 0, 1, \dots, E - 1, \quad (3)$$

with TE_0 the first echo time and ΔTE the spacing. Two field scales follow from ΔTE alone. The **wrap spacing** $1/\Delta \text{TE}$ separates neighbouring candidate solutions for the field. The **half-wrap** $1/(2\Delta \text{TE})$ is half that spacing, the distance from any candidate to the midpoint between it and its neighbour, and serves as the decision boundary for any rule that prefers the smallest field.

2.3 Recovering the offset

The construction below is MCPC-3D-S. Because θ is common to every echo, it cancels in a difference. For the first two echoes,

$$\Delta\varphi = W(\varphi_1 - \varphi_0) = W((\theta + 2\pi f t_1) - (\theta + 2\pi f t_0)) = W(2\pi f \Delta \text{TE}). \quad (4)$$

The middle step follows from Equation 2, since each φ_e differs from $\theta + 2\pi f t_e$ by whole turns. The wrap remains outside the subtraction because $\varphi_1 - \varphi_0$ is a difference of wrapped values.

$\Delta\varphi$ is itself wrapped, but a difference of neighbouring echoes varies slowly in space, so a region-growing algorithm can restore its missing turns. Spatial unwrapping recovers phase differences

between voxels rather than absolute turn counts, so what it returns may differ from the true unwrapped difference by a whole number of turns. Write $\Delta\varphi_{\text{uw}}$ for what it returns. Extrapolating to $t = 0$ then recovers the offset, with $k = \text{TE}_0/\Delta \text{TE}$,

$$\hat{\theta} = W(\varphi_0 - k \cdot \Delta\varphi_{\text{uw}}). \quad (5)$$

The outer W is required because $\Delta\varphi_{\text{uw}}$ is unbounded and $k \cdot \Delta\varphi_{\text{uw}}$ with it. That unrecovered number of turns is the only thing Equation 5 leaves open.

3 The candidate ladder

The turn count in $\Delta\varphi_{\text{uw}}$ is unknown, so consider every value it could take. Shifting what the unwrapper returned by whole turns,

$$\Delta\varphi_{\text{uw}}^{(N)} = \Delta\varphi_{\text{uw}} + 2\pi N, \quad N \in \mathbb{Z}, \quad (6)$$

produces one **candidate** per integer, with $N = 0$ the value the unwrapper (ROMEO) actually returned. A candidate is not a partial answer. Each one fixes an offset and a field together, and both follow from what has already been established.

The offset comes from putting Equation 6 into Equation 5. The shift enters multiplied by k , so

$$\hat{\theta}_N = W(\varphi_0 - k(\Delta\varphi_{\text{uw}} + 2\pi N)) = W(\hat{\theta}_0 - 2\pi k N). \quad (7)$$

The field is the slope of the unwrapped difference, phase gained per unit time over the interval ΔTE , so

$$\hat{f}_N = \frac{\Delta\varphi_{\text{uw}}^{(N)}}{2\pi\Delta \text{TE}} = \hat{f}_0 + \frac{N}{\Delta \text{TE}}. \quad (8)$$

One candidate is correct. Write N^* for its index, so that $\hat{\theta}_{N^*} \equiv \theta$ and $\hat{f}_{N^*} = f$. Nothing observable identifies N^* , which is the whole difficulty.

Two features of Equation 7 and Equation 8 matter later. The first is that a single integer fixes both quantities. There is no candidate carrying the correct field alongside an incorrect offset, or the reverse: they move together, the field in steps of $1/\Delta \text{TE}$ and the offset in steps of $2\pi k$. The second is that the candidates are **discrete**. The field is not uncertain within an interval; it is determined up to a set of isolated values spaced $1/\Delta \text{TE}$ apart, with nothing in between available as an answer. Whatever goes wrong later is therefore wrong by a whole wrap, never by a little.

3.1 The rule ROMEO applies

One candidate has to be chosen, and ROMEO's global correction is what chooses it. After spatially unwrapping, it subtracts 2π times the median rounded turn count over the mask, which leaves that median at zero,

$$\text{median } \lfloor \Delta\varphi_{\text{uw}}/2\pi \rfloor = 0. \quad (9)$$

That condition reads directly as a statement about the field. By Equation 8 a turn of the difference is $1/\Delta \text{TE}$ of field, so rounding the turn count to zero is the same as requiring

$$|\hat{f}_0| < \frac{1}{2\Delta \text{TE}}, \quad (10)$$

the field lying inside the **half-wrap**. Because neighbouring candidates are a full wrap apart, exactly one of them satisfies Equation 10, so the rule selects the candidate of smallest field magnitude: the rung of the ladder nearest zero. Equation 9 tests that condition through one particular statistic, an unweighted median of rounded turn counts over its own mask; which summary of the field is the right one is taken up in Section 7.

Equation 10 is an assumption rather than a measurement, and its justification is physical. Before imaging, the scanner sets its demodulation frequency to the centre of the water resonance, so the residual field is expected to be a small offset rather than an arbitrary one.

Where that expectation holds the rule returns N^* . Where it does not, it returns a neighbouring rung instead, and nothing in the data marks the difference.

Proposition 1. For evenly spaced echoes, every candidate predicts the same recorded phase. That is, $W(\hat{\theta}_N + 2\pi\hat{f}_N t_e)$ does not depend on N , at any echo.

The rule is therefore forced rather than merely convenient: no statistic computed from the phase can rank the candidates. Proved in Appendix A, which draws out the consequences and sets out what the result does not cover.

4 Why the rule fails on some frames

A wrong candidate, chosen once and applied to a whole acquisition, would matter very little. Every frame would be displaced by the same $1/\Delta$ TE of field, and the field offset is already referenced to an arbitrary demodulation frequency, so a fixed additive constant changes neither temporal contrast nor spatial structure.

The difficulty is that the choice is not made once. Three properties of Equation 9 combine to produce a failure.

The first is that it is evaluated on every frame. A functional acquisition is processed frame by frame, and each frame applies Equation 9 to its own data, without reference to the frames around it.

The second is that the left side of Equation 9 is a sample median over a mask of voxels, each carrying its own field. It absorbs whatever varies between frames, including physiological change, subject motion, and noise in the voxels the mask admits.

The third is that Equation 10 is a comparison against a fixed boundary. When $|\hat{f}_0|$ is far below $1/(2\Delta \text{ TE})$ the margin is wide and no plausible perturbation of the median changes which candidate satisfies the inequality. As the field approaches the half-wrap that margin approaches zero.

The practical consequence is that ROMEO's global correction is occasionally off by a single turn. For a subject whose field lies near the half-wrap the candidate is chosen on each frame independently, by a noisy median sitting close to a threshold, so different frames fall on different sides of it. Because the candidates are a full wrap apart, a frame that falls on the wrong side is displaced by $1/\Delta$ TE rather than by a small amount. The resulting time series is stable to a fraction of a hertz across most frames and displaced by a full wrap on a scattered subset, with no intermediate values.

The damage is temporal. A uniformly wrong field is harmless, but a field taking two values within a single acquisition corrupts the temporal structure the acquisition was performed to measure.

5 Detecting a failed choice

By Prop. 1 no test can rank one candidate against another. The test in this section asks a different question: whether what the pipeline has produced is a candidate at all.

The question has content because the offset and the field come from two separate steps. The first uses $\Delta\varphi_{\text{uw}}$ to form the offset, by Equation 7. The second removes that offset from every echo, spatially unwraps the first echo — applying Equation 9 a second time, now to offset-removed phase — and unwraps the remaining echoes against it to recover the field. Nothing carries the first step’s choice of shift into the second.

Consider a frame on which the first unwrapping has tipped, so the offset removed belongs to a neighbouring rung rather than the correct one. The echoes still carry the correct field, and the second unwrapping recovers it, because Equation 9 prefers the small field and in offset-removed data the correct field is the small one. What comes out therefore has an offset from one rung and a field from another. That combination is not a rung of the ladder, so Prop. 1 does not cover it, and it leaves a signature in the data.

5.1 The signature

Removing an offset $\hat{\theta}_N$ from the recorded phase leaves

$$\psi_e = \varphi_e - \hat{\theta}_N, \quad (11)$$

a difference of two wrapped values, and therefore not itself a phase. Unwrapping it in the echo direction and fitting a line,

$$\psi_e \approx c + s \cdot t_e, \quad (12)$$

gives a slope s , which is the field, and an intercept c , the value the line extrapolates to at $t = 0$.

The intercept is the signature. Removing the offset that belongs to a given field leaves phase that starts at zero, so a matched offset and field give a line through the origin, $c = 0$. A mismatched pair does not, and by how much is fixed.

Proposition 2. Let the echo unwrap return the correct field, while the offset removed was $\hat{\theta}_N$ for some $N \neq N^*$. Then

$$\psi_e \equiv 2\pi f t_e + c \pmod{2\pi}, \quad c = W(2\pi k(N - N^*)), \quad (13)$$

with c identical at every echo.

Proof. By Equation 7, $\hat{\theta}_N \equiv \hat{\theta}_{N^*} - 2\pi k(N - N^*) \equiv \theta - 2\pi k(N - N^*)$. Then $\psi_e \equiv (\theta + 2\pi f t_e) - (\theta - 2\pi k(N - N^*)) = 2\pi f t_e + 2\pi k(N - N^*)$, and the constant does not depend on e . ■

Two properties make c usable in practice. Being identical at every echo, it displaces the whole fitted line rather than averaging away across echoes. Being one global constant rather than a per-voxel effect, it can be aggregated over a mask without dilution.

5.2 What the intercept can and cannot settle

The intercept measures neither an error in the offset nor an error in the field. It measures whether the two are mutually consistent, and consistency between them is the one property the data can check.

That makes it a filter rather than a ranking. Among the shifts of Equation 6, those with a large intercept are not candidates and can be discarded. Those with $c \approx 0$ are candidates, and Prop. 1

then applies in full: they predict identical data and the intercept cannot separate them. Choosing among the survivors still requires Equation 10.

Both steps are needed, for different reasons. The filter is needed because a shift can report a reasonable-looking field while carrying an offset that does not generate it, and since that offset is subsequently removed from every echo, the error propagates into everything downstream. The rule is needed because the filter leaves more than one survivor. What the filter contributes is that when the first unwrapping tips, the surviving set is not the same set, and applying the rule to the shifted set recovers the frame.

6 The selection algorithm

The procedure runs once per frame and returns one shift N to apply to $\Delta\varphi_{\text{uw}}$.

Step 1. Consider $N \in \{-1, 0, +1\}$. One step is a full wrap of global field, and Equation 9 has already brought the field inside the half-wrap, so a correction of more than one wrap is not reachable.

Step 2. For each N , form $\hat{\theta}_N$ by Equation 7, remove it from both echoes, unwrap in time, fit Equation 12, and record $|c|$ aggregated over a conservative brain mask.

Step 3. Discard any candidate whose $|c|$ exceeds a cutoff. By Prop. 2 the candidates of Step 1 have intercepts 0 and $\pm W(2\pi k)$, so the cutoff is placed at half that separation, making the test a nearest-neighbour assignment between two possibilities rather than a tuned threshold. The implementation also measures the largest intercept observed among the candidates and adopts the stricter of the two scales, which keeps the test honest when the echo unwrap has already absorbed part of the residue.

Step 4. If one candidate survives, take it, whatever its field. If none survives, change nothing, since a failed fit is not evidence for any candidate. If several survive, they are indistinguishable to the data, and Equation 10 decides among them by smallest $|\tilde{f}_N|$, the field the reconstruction actually returned.

Steps 3 and 4 answer different questions and neither suffices alone. Step 3 uses evidence, rejecting candidates the data contradicts. Step 4 uses the assumption of Equation 10, and applies it only among candidates the data has already accepted, which is what stops the assumption from overriding the evidence.

Step 4 restates Equation 10, so it is worth being clear about what it changes. Equation 9 evaluates the condition as an unweighted median of **rounded** turn counts over a dilated mask. Rounding discards the sub-wrap information that says how close the decision was, and dilation admits edge voxels whose fields are extreme. Step 4 evaluates $|\tilde{f}_N|$ directly, weighted, over a conservative mask. The condition is the same; the statistic is not, and Section 7 is about why that matters.

One degenerate case deserves mention. If k is a whole number then $W(2\pi k(N - N^*)) = 0$ for every N , so the intercept carries no information. By Equation 7 the offset $\hat{\theta}_N = W(\hat{\theta}_0)$ is then also independent of N , so no choice alters the output, and leaving $N = 0$ is correct. A guard on $\delta = 0$ states that intent, but it need not be what enforces it: in floating point $W(2\pi k)$ evaluates to a rounding residue near 10^{-16} rather than to zero, and the cutoff that residue sets is then small enough that every candidate fails it, so the empty-survivor case of Step 4 returns $N = 0$ regardless. Either route gives the same answer.

7 Estimating the field

Step 4 compares $|\tilde{f}|$ against the half-wrap, so the reliability of the comparison depends on what that estimate measures and on how steady it is between frames.

Equation 10 is a statement about **bulk tissue**, since bulk tissue is what the scanner’s frequency adjustment centres. It is not a statement about the mean field, nor about the median field over whatever mask happens to be available.

A median is used rather than a mean or an M-estimator. The bulk distribution is broad and right-skewed, with long tails from air-tissue interfaces, sinuses and mask edges, and the half-wrap boundary falls within its body rather than out in a tail. A statistic pulled toward the mean therefore reads systematically closer to the boundary without describing bulk tissue any more accurately.

MEDIC uses that median over an eroded brain mask, weighted by squared magnitude. Both choices work against the tails: erosion removes edge voxels geometrically, and the weighting suppresses low-signal voxels, which the tails over-represent. The weight is taken from the second echo rather than the first, because the field is read from a difference whose noise is dominated by the weaker of the two.

8 The algorithm in full

Collected here as a specification. All quantities are per-voxel unless a reduction over a mask is written explicitly, and the whole procedure is applied independently to each frame.

Input. Magnitudes m_0, m_1 and wrapped phases φ_0, φ_1 at the first two echo times TE_0 and TE_1 ; an unwrapping mask Ω ; a conservative mask $\Omega_c \subseteq \Omega$ for reductions.

Output. The phase offset $\hat{\theta}$.

Constants. $\Delta \text{TE} = \text{TE}_1 - \text{TE}_0$, $k = \text{TE}_0 / \Delta \text{TE}$, $\delta = |W(2\pi k)|$.

1. Unwrap the phase difference.

$$S = m_0 m_1 e^{i(\varphi_1 - \varphi_0)}, \quad \Delta\varphi = \arg S, \quad \Delta\varphi_{\text{uw}} = \text{ROMEO}(\Delta\varphi; |S|, \Omega). \quad (14)$$

ROMEO returns a spatially unwrapped field with Equation 9 imposed.

2. Reconstruct each candidate. For $N \in \{-1, 0, +1\}$:

$$\begin{aligned} \hat{\theta}_N &= W(\varphi_0 - k(\Delta\varphi_{\text{uw}} + 2\pi N)), \\ (\psi_0^{(N)}, \psi_1^{(N)}) &= \text{ROMEO}_{4D}((\varphi_0, \varphi_1) - \hat{\theta}_N; (m_0, m_1), \Omega, (\text{TE}_0, \text{TE}_1)), \\ s^{(N)} &= (\psi_1^{(N)} - \psi_0^{(N)}) / \Delta \text{TE}, \quad \tilde{f}_N = s^{(N)} / 2\pi, \\ c_N &= \left| \text{med}_{\Omega_c} [\psi_0^{(N)} - s^{(N)} \text{TE}_0] \right|. \end{aligned} \quad (15)$$

ROMEO_{4D} imposes Equation 9 a second time, which is why $\tilde{f}_N$ need not equal the ladder field $\hat{f}_N$ of Equation 8, and so why c_N can be nonzero at all.

3. Set the rejection scale.

$$\sigma = \min\left(\max_N c_N, \delta\right), \quad \mathcal{C} = \{N : c_N < \sigma/2\}. \quad (16)$$

If δ is zero, return $\hat{\theta}_0$: k is then integral, so by Equation 7 the offset does not depend on N and no choice alters the output. Test this within a tolerance rather than against exact zero, since $W(2\pi k)$ evaluates to a small rounding residue in floating point.

4. Select.

$$\hat{N} = \begin{cases} 0 & \text{if } \mathcal{C} = \emptyset \\ N & \text{if } \mathcal{C} = \{N\} \\ \arg \min_{N \in \mathcal{C}} |\text{med}_{\Omega_c}^w \tilde{f}_N| & \text{otherwise, } w = m_1^2. \end{cases} \quad (17)$$

5. Return $\hat{\theta}_{\hat{N}}$.

Three details of the specification carry the reasoning of the earlier sections. The two reductions differ deliberately: step 2 takes a plain median, since it estimates one global constant, while step 4 takes a median weighted by m_1^2 of a distribution whose shape matters, for the reasons in Section 7. The scale σ in step 3 takes the smaller of the analytic separation δ and the largest intercept observed, which makes the consistency test harder to pass and so biases the procedure toward leaving $\hat{N} = 0$. And the cases in step 4 are ordered, not interchangeable: an empty $\mathcal{C}$ returns 0 rather than consulting Equation 10, a singleton is taken whatever its field, and only a tie reaches Equation 10.

8.1 What the components must provide

Three things are deliberately left open: the spatial unwrapper, the two masks, and the weighted median. What matters is the property each must satisfy.

The unwrapper. Steps 1 and 2 need one that imposes Equation 9 on what it returns, step 2 on the echo it unwraps first. Nothing else about it is load-bearing: the selection logic reads only the returned phase, and Section 5 needs only that the two unwrapping steps choose their shifts independently, since that is what pairs an offset from one rung with a field from another and leaves a signature to detect.

The masks. Ω is generous, brain plus a margin, because the unwrapper must grow through the whole region of interest. Ω_c is interior, excluding edge and low-signal voxels, because both reductions summarise a distribution whose tails come from exactly those voxels (Section 7). Nothing depends on how either is built, only on their differing in that way, with Ω_c well inside Ω .

The weighted median. med^w is the value at which cumulative weight, over voxels sorted by value, first reaches half the total. The convention matters only at an exact tie, where any deterministic choice will do.

8.2 Reference data

The failure of Section 4 can be reproduced on public data. [ds006131](#) (v2.0.1, “PAFIN”) acquires five echoes at $\text{TE}_0 = 14.2$ ms and $\Delta \text{TE} = 24.73$ ms, so $\delta = 2.68$ rad and the half-wrap is 20.22 Hz. Two of its 51 subjects carry a bulk field close enough to that boundary for Equation 9 to tip on individual frames:

ds006131/sub-<id>/ses-1/func/

sub-<id>_ses-1_task-bao_dir-AP_run-01_echo-{1..5}_part-{mag,phase}_bold.nii.gz

$\hat{N} \neq 0$ on 19 of 243 frames for sub-20828 and on 158 of 243 for sub-24630.

Afterwards neither field-map time series contains a wrap-sized step, where before each carried one on every affected frame. The counts depend on the masks in use, so treat them as a guide: what is robust is that these two subjects need a nonzero $\hat{N}$ on a scattered subset of frames and the other 49 do not. An implementation returning $\hat{N} = 0$ throughout on them has not reproduced the correction; one returning nonzero shifts across many subjects is over-correcting.

A Proof of Proposition 1

Proof of Prop. 1. Working modulo 2π , so that W of either side agrees, Equation 7 and Equation 8 give

$$\begin{aligned}\hat{\theta}_N + 2\pi\hat{f}_N t_e &\equiv \hat{\theta}_0 - 2\pi kN + 2\pi(\hat{f}_0 + N/\Delta \text{ TE})t_e \\ &= \hat{\theta}_0 + 2\pi\hat{f}_0 t_e + 2\pi N(t_e/\Delta \text{ TE} - k).\end{aligned}\tag{18}$$

By Equation 3, $t_e/\Delta \text{ TE} = k + e$, so the parenthesis is the integer e and the final term is $2\pi Ne$, a whole number of turns. Both sides therefore wrap alike. ■

A.1 The geometry




Figure 1: Candidates as lines in phase against echo time, drawn for $k = 0.6$. All of them meet at TE_0 , so the offset step $2\pi k$ is exactly the phase a one-wrap steeper slope gains by the first echo. Past TE_0 the lines separate by one further turn per echo spacing, so at echo e they stand whole turns apart (open circles) and wrap to identical samples.

Figure 1 is the same statement drawn. Against echo time each candidate is a straight line, intercept $\hat{\theta}_N$ and slope $2\pi\hat{f}_N$. Setting $t = \text{TE}_0$ in the last line of the proof kills the parenthesis, so every candidate takes the value $\hat{\theta}_0 + 2\pi\hat{f}_0 \text{TE}_0$ at the first echo: the candidates form a pencil of lines through one point. Past that point they separate at one turn per echo spacing, standing $2\pi Ne$ apart at t_e . Wrapped phase cannot see a whole turn, so every line reproduces the same measured samples.

The two step sizes are consequences rather than choices. Requiring two candidates to differ by whole turns at every echo forces their slopes to differ by one turn per echo spacing, which is the field step $1/\Delta \text{ TE}$; that fixed, their offsets must differ by the $2\pi k$ such a slope has already gained at TE_0 , which is exactly what makes the two lines meet there.

A.2 Consequences

Three consequences follow, and the rest of the note rests on all of them.

The rule is unavoidable. Since the candidates predict identical data, no statistic computed from the phase can rank them. Equation 10 is the only kind of answer available, and any implementation must supply something like it.

A longer echo train does not help. The cancellation holds at every echo at once, so a five-echo acquisition constrains N no better than a two-echo one. The later echoes are not independent measurements of the ambiguity. This depends on the even spacing of Equation 3 rather than on phase imaging in general, but an EPI train is evenly spaced by construction.

The undetectable errors are exactly enumerated. Prop. 1 applies to a candidate as defined in Equation 7 and Equation 8, an offset paired with the field of the same N . Nothing outside that set is protected, which is what makes any detection possible at all, and Section 5 is about exploiting it.