Supplementary

0.1 Inclusion Model and Effective Susceptibility

Magnetic field perturbations Susceptibility contrast between phases ($\Delta\chi$) gives rise to a dipolar magnetic field perturbation B_d, together with a Lorentz cavity (local-field) contribution B_L associated with the surrounding medium. The resulting total field perturbation ΔB generated by a single spherical inclusion, shown in Fig. 1.a, is expressed in a spherical coordinate system as:

$$\Delta B = B_\mathrm{d} + B_\mathrm{L} = \frac{\Delta\chi}{3}\frac{R^3}{\rho^3}(3\cos^2\theta - 1)B_0 + \frac{1}{3}\chi_m B_0$$

where θ and ρ are the spatial coordinates, R is the radius of the inclusion, and $\Delta\chi = \chi_i - \chi_m$ is the susceptibility contrast, where the subscripts i and m denote the inclusion and matrix phases, respectively. In a macroscopically isotropic medium, the average field contribution over a volume V can be computed as:

$$\overline{\Delta B} = \frac{1}{V}\int_V B_\mathrm{d}\mathrm{d}V + \frac{1}{V}\int_V B_\mathrm{L}\mathrm{d}V$$

The first integral can be evaluated explicitly in spherical coordinates:

$$\overline{B_\mathrm{d}} = \frac{\Delta\chi R^3 B_0}{3V}\int_0^{2\pi}\int_R^{R_V}\int_0^{\pi}\frac{1}{\rho^3}(3\cos^2\theta - 1)\rho^2\sin\theta\mathrm{d}\theta\mathrm{d}\rho\mathrm{d}\varphi = 0.$$

Physically, this result indicates that the dipole field does not introduce a net magnetic field contribution when averaged over space, but rather redistributes the existing magnetic field locally around the inclusion. The average Lorentz contribution becomes:

$$\overline{B_\mathrm{L}} = \frac{B_0}{3V}\int_V \chi(\mathbf{r})\mathrm{d}V, \tag{3}$$

where $\chi(\mathbf{r})$ denotes the spatial susceptibility distribution. To bridge this single-inclusion description with the macroscopic response of a dense mixture, the local field perturbations can be integrated into an effective-medium framework. In this work, we model the composite medium using a host-referenced effective susceptibility χ_eff such that $\overline{B_\mathrm{L}} = B_0(\chi_\mathrm{eff} - \chi_m)/3$.

Formulation of a effective medium model in MRI-relevant susceptibility space The present formulation recasts the classical Maxwell Garnett (MG) effective-medium framework in terms of host-referenced susceptibility contrasts, consistent with what measured through MRI measurements.

Standard MG theory is a mean-field effective-medium model in which: i) each inclusion is represented by α, i.e. its dipolar response to a local field H_{loc}, and ii) H_{loc} is related to the macroscopic field H through a Lorentz cavity model. In classical electromagnetism, this framework is typically formulated in terms of absolute constitutive parameters, so α is expressed in terms of magnetic permeability μ, and Lorentz cavity model for self-consistency is imposed on the magnetization M. Nonetheless, the measurable signal in MRI is inherently background-referenced: observable phase and frequency shifts arise from spatial field inhomogeneities arising from susceptibility contrasts rather than from absolute magnetization. This motivates a MG-type closure written directly in terms of host-referenced contrast variables. In the present formulation, the MG mean-field *structure* (dipolar response combined with a Lorentz cavity model) is retained, but is applied on MRI-relevant contrast quantities.

The presented model is then characterized by two explicit modelling choices, motivated by the fact that MRI measurements are relative to χ_m, and driven by $\Delta\chi$–induced dipolar field perturbations (ΔM). Specifically, we: i) parametrize the single-inclusion dipolar response using the susceptibility-based factor α_χ rather than the classical μ-based form, and ii) we we apply the Lorentz cavity model to the host-normalized contrast variable $\mathcal{X} = \Delta M/\chi_m$ rather than the absolute magnetization M. These choices are made since, by formulating the derivation selecting α_χ as the inclusion response and $\mathcal{X}$ as the closure variable for the Lorentz cavity model, the resulting model in the same variable space as the one accessible through MRI measurements.

Derivation of χ_eff In this model, the host-normalized contrast state variable entering the Lorentz cavity model is defined as:

$$\mathcal{X} \equiv \frac{\Delta M}{\chi_m},$$

where ΔM is the excess magnetization relative to the host $\Delta M \equiv M_\mathrm{tot} - M_m$. In terms of the effective susceptibility, the contrast magnetization satisfies

$$\frac{\Delta M}{H} = \chi_\mathrm{eff} - \chi_m. \tag{4}$$

Here χ_{eff} represents an MRI-relevant effective susceptibility rather than a strictly homogenized constitutive parameter in the classical effective-medium sense.

Since the frequency shifts measured in MRI experiments are driven by susceptibility-induced field perturbations, we parameterize the dipolar response of spherical inclusions using a susceptibility-based Clausius–Mossotti-type factor:

$$\alpha_\chi = \frac{\chi_i - \chi_m}{\chi_i + 2\chi_m}. \tag{5}$$

While the standard MG sphere factor is commonly expressed in terms of magnetic permeabilities (i.e., α_μ), we adopt the susceptibility-space form in Eq.5 as an MRI-relevant contrast parameterization. In this sense, α_χ encodes the modeling choice that the inclusion response relevant to MRI is governed by susceptibility contrast rather than absolute permeabilities.

In the MG mean-field approximation, the magnetization is represented as a dipole density driven by the local field. Framing the same concept in terms of our state variable $\mathcal{X}$, we write:

$$\mathcal{X} = 3\phi\,\alpha_\chi\,H_{\text{loc}}, \tag{6}$$

Next, we introduce the Lorentz-closure. Here, the mean-field self-consistency is imposed on the contrast through $\mathcal{X}$. This step is done assuming the Lorentz spherical geometric factor (1/3) can be retained. The cavity model then becomes:

$$H_{\text{loc}} = H + \frac{\mathcal{X}}{3} \tag{7}$$

This retains the Lorentz geometric structure while shifting the self-consistency condition to a contrast variable appropriate for MRI. Combining Eq.6 and 7,

$$\mathcal{X} = 3\phi\,\alpha_\chi\left(H + \frac{\mathcal{X}}{3}\right) \rightarrow \frac{\mathcal{X}}{H} = \frac{3\phi\,\alpha_\chi}{1-\phi\alpha_\chi}.$$

Mapping back to magnetization contrast via $\Delta M = \chi_m \mathcal{X}$ yields

$$\frac{\Delta M}{H} = \chi_m\,\frac{3\phi\,\alpha_\chi}{1-\phi\alpha_\chi}.$$

Using Eq.4, the host-referenced effective susceptibility becomes

$$\chi_{\text{eff}} - \chi_m = \chi_m\,\frac{3\phi\,\alpha_\chi}{1-\phi\alpha_\chi}.$$

Substituting Eq.5 yields the closed-form expression

$$\chi_{\text{eff}} - \chi_m = \frac{3\chi_m\,\phi\,\Delta\chi}{3\chi_m + \Delta\chi(1-\phi)}. \tag{8}$$

Eq. 8 defines an MG-type self-consistent closure written in terms of host-referenced susceptibility contrast. The formulation preserves the MG/Lorentz mean-field *structure* (dipole-density response and spherical Lorentz geometric factor), while adopting two MRI-relevant modelling choices, and the resulting expression defines an effective description of background-referenced field perturbations relevant to MRI. Additional remarks on Eq. 8 are reported at the end of this section.

Frequency shift definition Accordingly, Eq. 3 is modelled using the host-referenced effective susceptibility given by Eq. 8, yielding:

$$\overline{\Delta B} = \frac{B_0(\chi_{\text{eff}} - \chi_m)}{3} \tag{9}$$

Finally, the susceptibility-induced frequency shift (expressed in parts-per-million, ppm) is defined as:

$$\delta_f = \frac{\tilde{\omega}_0 - \omega_0}{\omega_0} 10^6 \tag{10}$$

where $\omega_0 = \gamma B_0$ is the unperturbed Larmor frequency of water and $\tilde{\omega}_0 = \gamma(B_0 + \overline{\Delta B})$ is the perturbed Larmor frequency of the heterogeneous system.

Combining Eq. 9 and Eq. 10 yields:

$$\delta_f = \frac{\chi_{\text{eff}} - \chi_m}{3} 10^6 = \frac{\chi_m\,\phi\,\Delta\chi}{3\chi_m + \Delta\chi(1-\phi)} 10^6.$$

Remarks We hereby report some remarks on Eq.8.

i) **Dilute vs dense limit.** In the dilute limit ($\phi \to 0$) inclusions barely affect the system, thus $\chi_{\text{eff}} \to \chi_m$. Conversely, when $\phi \to 1$ the system transitions to an inclusion dominated susceptibility: $\chi_{\text{eff}} \to \chi_i$.

ii) **Linearity in the dilute regime.** As $\phi \to 0$ the response is linear in ϕ in the dilute regime, consistent with the superposition of non-interacting dipoles:

$$\chi_{\text{eff}} - \chi_m \approx \frac{3\chi_m \Delta\chi}{3\chi_m + \Delta\chi}\phi. \tag{11}$$

Furthermore, in the limit of small contrast ($\Delta\chi \to 0$) $\chi_{\text{eff}} - \chi_m \approx \Delta\chi \cdot \phi$, as in simplified linear mixing.

iii) **Sign and monotonicity.** Provided the denominator in Eq.8 remains positive, the sign of $\chi_{\text{eff}} - \chi_m$ matches the sign of $\Delta\chi$, and the magnitude increases with ϕ, consistent with the mixture moving toward the inclusion property.

iv) **Self-consistency and nonlinearity.** The nonlinear (rational) dependence on ϕ arises from the Lorentz cavity model, which incorporates mean-field interaction corrections from dipole-dipole coupling beyond the linear dilute approximation.

v) **Host/inclusion asymmetry.** As standard MG, this closure is host-based and not symmetric under swapping phases. For the specific form in Eq.8, the limit $\phi \to 1$ yields $\chi_{\text{eff}} \to \chi_i$. However, MRI signal comes from water molecules in the system, thus decreasing as $\phi \to 1$. The asymmetry of this model is therefore appropriate as it treats inclusions as dispersed in a continuous host phase, as opposed to symmetric models (e.g., Bruggeman) in which both phases are equally represented.

vi) **Geometry generalization.** The factor 1/3 is the Lorentz factor for spherical cavities. Nonetheless, the same mean-field structure can be extended to non-spherical inclusions by replacing 1/3 with an appropriate demagnetization factor, preserving the closure form while modifying the geometric prefactor.

0.2 Volume-Averaged Volume Fraction Measurements Through FIDs in Bubbly Flows

Free Induction Decays (FIDs) are considered the simplest MRI pulse sequence, generally consisting of a single hard RF pulse, followed by the read gradient. This sequence is commonly used provide volume-averaged information about the relaxation time and water content (in H^1 NMR) within the coil sensitive volume. This property is particularly useful in presence of high susceptibility contrast systems, as the large $\Delta\chi$ prevents the use of amplitude images from being a reliable validation of the volume-averaged volume fraction ($\overline{\phi}$). Thus, in this study we use a standard FID sequence to measure $\overline{\phi}$ during the bubbly bubbly flow experiments.

FIDs measure the spin relaxation toward the z-axis over time, producing a exponential decay signal $A(t)$:

$$A(t) = A_0 e^{-TE/T_2^*}, \tag{12}$$

where A_0 is the signal amplitude at $t = 0$ and is proportional to the water content within the coil sensitive volume. The echo-time TE is the time between the delivery of the radio frequency pulse and the peak of the received signal. The decay constant T_2^* represents the effective transverse relaxation time, which is determined by both spin–spin interactions and magnetic field inhomogeneities.

When inclusions (thus $\Delta\chi$) are present, two effects arise: i) the presence of the inclusions reduces the volume of the liquid proportionally to ϕ, leading to a reduction of A_0, and ii) inclusions lead to a reduction of T_2^*, which arise due to the field inhomogeneities induced by the presence of a $\Delta\chi$. Specifically, the total spin-spin relaxation time of the system follows [86]:

$$\frac{1}{T_2^*} = \frac{1}{T_2} + C\Delta\chi B_0, \tag{13}$$

with T_2 the true relaxation time, only induced by the spin-spin interactions, and the second term accounts for the susceptibility differences, in which C a constant depending on the bubbles size and spatial distribution.

Figure 5.a shows the time domain FID signals measured for different ϕ, in which both these effects are visible. The reduction of T_2^* is reflected in a faster decay, as expected from Eq. 12, whereas the decrease of water content in the coil sensitive region leads to a lower A_0 ($t \to 0$). Thus, the signal decay can be expressed as an unknown function of time and volume fraction: $A(t, \phi)$. Denoting $A_\phi = A(0, \phi)$ as the FID initial amplitude in the presence of bubbly flow, and $A_{ref} = A(0, 0)$ as the same quantity when no air is injected, we define:

$$\delta_A = \frac{|A_\phi - A_{ref}|}{A_{ref}}, \tag{14}$$

where δ_A represents the percentage change in the initial FID amplitude due to the introduction of air bubbles. Low values of δ_A indicate a small change relative to the initial water content, corresponding to a low volume

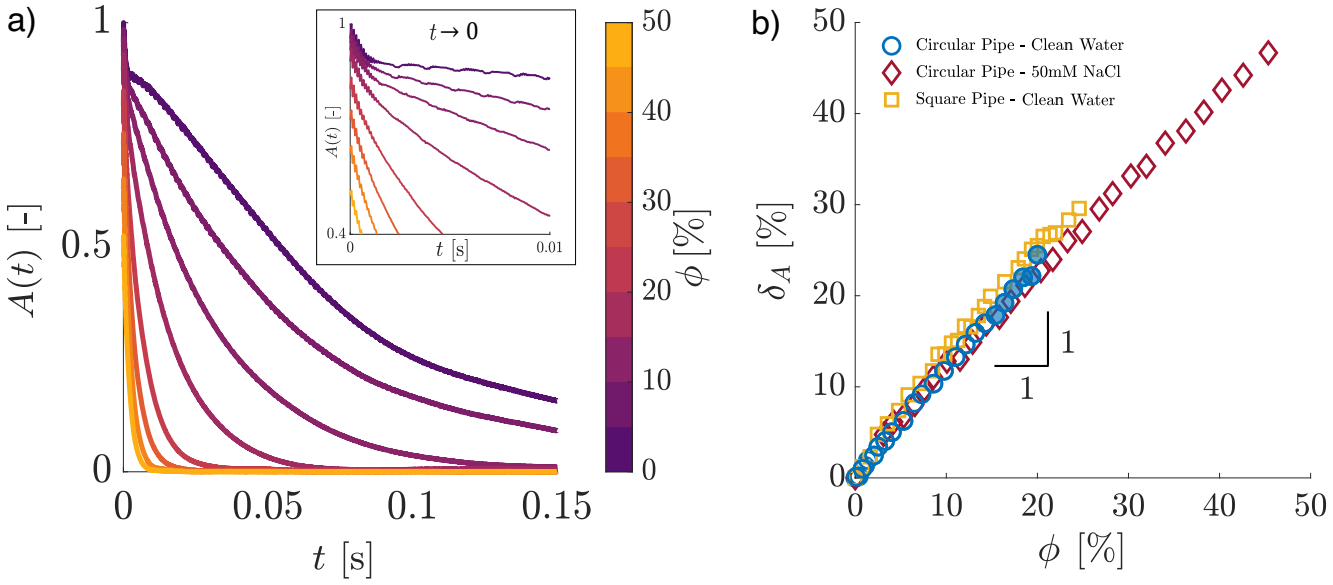

Figure 5: **a** Time domain FID for different volume fractions, with detail on reduction of A_0 ($t \to 0$). **b** Initial amplitude variation for different systems as a function of the volume fraction. Filled blue markers denote flows where bubble coalescence is observed.

fraction. Whereas, δ_A increases for higher volume fraction, as the water content reduces. Figure 5.b shows δ_A for the experiments conducted in both a square and a circular pipe. The bulk void fraction $\overline{\phi}$ is determined as $\overline{\phi} = \Delta h/(\Delta h + h_0)$, being h_0 the initial water height in the column and Δh the increase in the water level when air is injected. For the circular pipe, data are reported for both clean water and a 50 mM NaCl–water solution, in which bubble coalescence is inhibited. In the square pipe, the maximum measured volume fraction is 25% due to technical limitations related to the column height. For the clean water experiments the maximum volume fraction is 20% as excessive bubble coalescence resulted in the transition to a slug flow. When observed, bubble coalescence is highlighted by filling the markers. These results show a linear relationship between δ_A and the volume fraction, with a unitary slope. This is regardless of the system geometry and the presence or absence of bubble coalescence. This behaviour is expected, as A_ϕ is directly proportional to the water content of the system only. Consequently, δ_A is by definition the air-to-water ratio, which is equivalent to ϕ. Thus, this method allows direct measurement of $\overline{\phi}$ in the coil sensitive volume without requiring any additional calibration.

0.3 MRI Scanner Characterization

To characterize the magnet, we investigate: i) the effect of preheating the gradient coils before the acquisitions, ii) the steady-state stability of the main magnetic field over time, and iii) the effect of progressive averaging on the measurement noise.

0.3.1 Magnet Preheating

During operation, the gradient coils heat up, producing an additional temperature-induced frequency shift. This contribution must be separated from susceptibility-induced frequency shifts to ensure accurate measurements. To isolate the thermal effect, experiments are performed using a pure water sample, such that the observed frequency evolution arises mainly from coil heating. To assess whether intra-acquisition cooling affects the measurements, experiments were conducted with varying numbers of averages per acquisition. Figure 6 shows the measured frequency shifts for two different acquisition routines. This is done as between acquisitions the gradient coils remain inactive for several seconds, whereas during averaging within a single acquisition they operate continuously. The data are normalized by the maximum shift observed in each experiment: although the absolute offsets are substantial (typically on the order of hundreds of ppb) and therefore must be accounted for, their specific values depend strongly on the initial coil temperature and other experiment-dependent conditions. Consequently, the present analysis focuses on the temporal evolution of the heating process rather than on the absolute magnitude of the shifts.

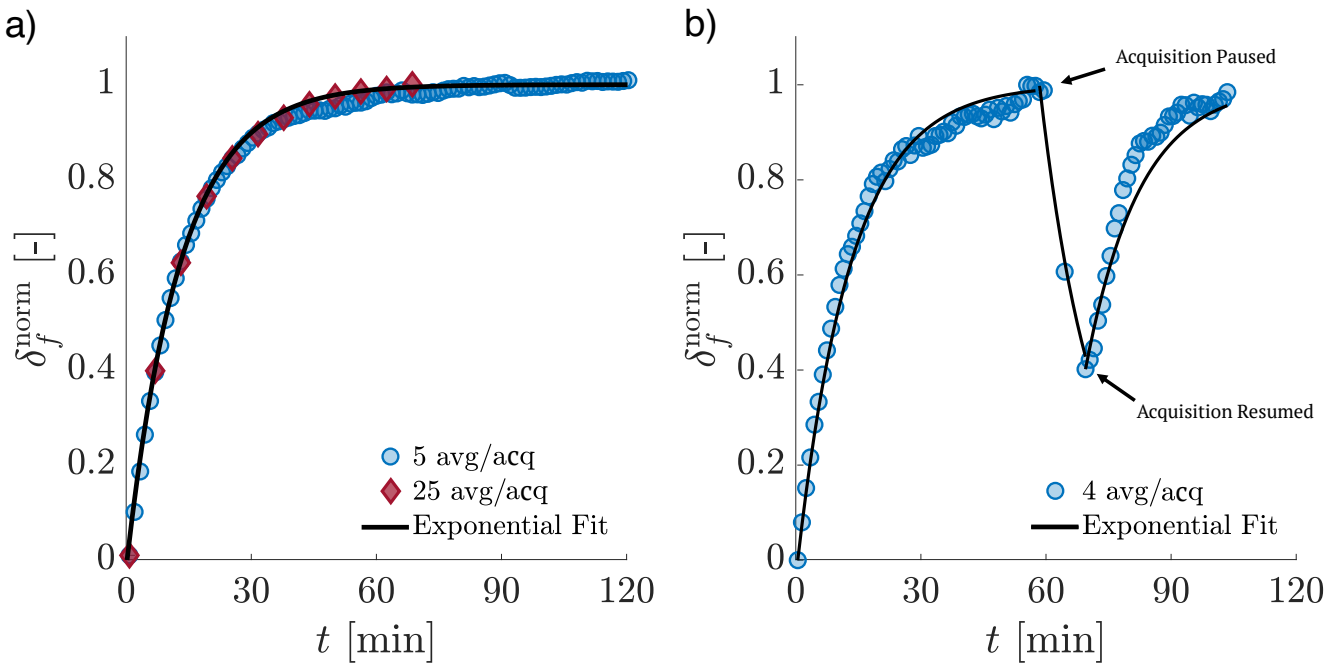

Figure 6: **a** Normalized temperature-induced frequency shift for two different PC-MRI acquisition routines. The fitted exponential saturation profile is superimposed. **b** Normalized temperature-induced frequency shift for intermittent acquisitions, in which coils inactivity leads to a reduction in the shift due to intra-acquisition cooling.

As shown in Fig. 6a, both acquisition schemes exhibit identical exponential heating behaviour $\delta_f^{\text{norm}} = 1 - e^{t/\tau}$, characterized by a time constant $\tau = 13$ min. Assuming that thermal steady-state is reached after approximately $t \approx 5\tau$, the system requires about one hour of preheating under the present experimental conditions. This requirement is independent of the number of averages per acquisition and is necessary to suppress temperature-induced frequency drifts. In Fig. 6b the same experiment is done also in case of intermittent acquisition, common when different samples are imaged and therefore need to be introduced and removed from the magnet.

0.3.2 Stability of the Main Magnetic Field

Once thermal steady-state is reached, fluctuations of the main magnetic field B_0 become a non-negligible source of frequency shift. These fluctuations are on the order of tenths of ppb and are therefore comparable to susceptibility-induced shifts in low-$\Delta\chi$ systems. Such variations are intrinsic to the instrument and arise from temporal instabilities of the main magnetic field. Because they cannot be eliminated, they define the fundamental uncertainty for susceptibility-based measurements. A pure water sample is used to quantify this drift at steady state.

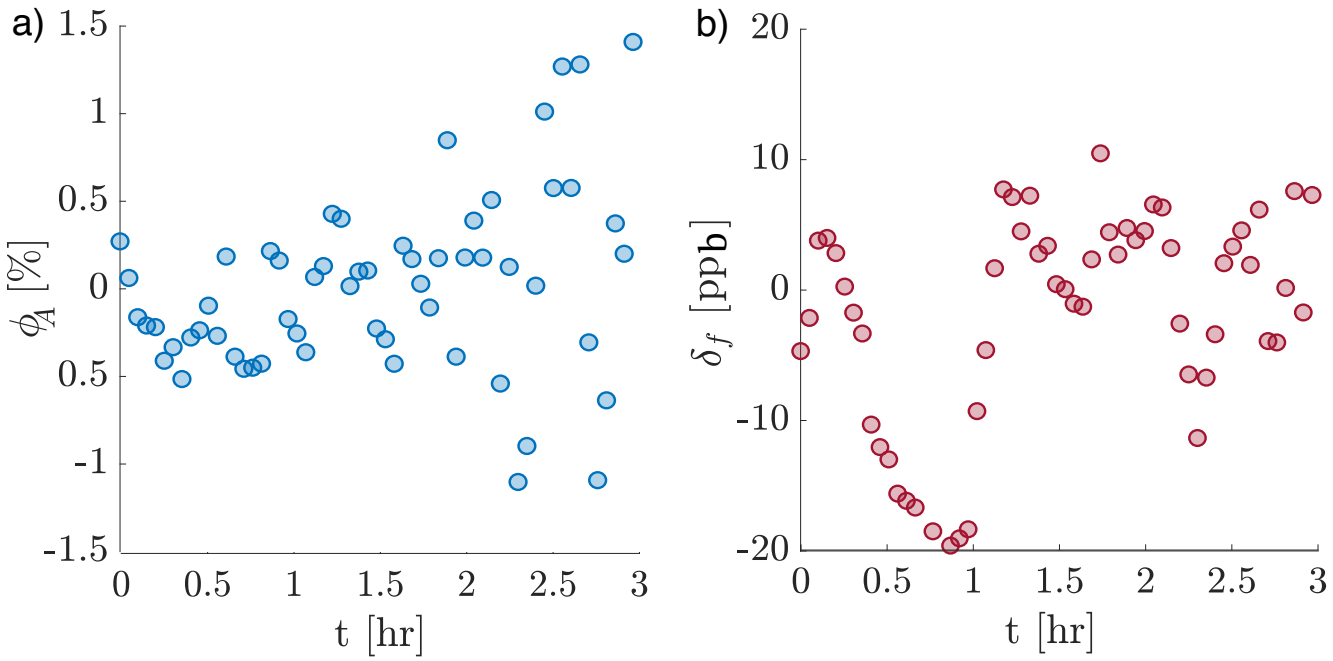

Figure 7: Time evolution of the amplitude-based volume fraction (**a**) and of the frequency shift (**b**) for a pure water sample. Both shifts are caused by the temporal drift of the main magnetic field and lead to an uncertainty in the measurements of both amplitude- and phase- based volume fractions.

Following preheating and subsequent shimming, the system is operated continuously using the same sequence as in the main experiments for amplitude and phase acquisition. From these data, the volume fraction derived from amplitude images (ϕ_A) and the frequency shift (δ_f) are measured over time. The data present a slow temporal drift, which is first removed from the measured signal, and the uncertainty associated with B_0 stability is then quantified as the standard deviation of the remaining fluctuations. The two drift-compensated time series are shown in Fig. 7a and b, respectively. This analysis yields an uncertainty of $\pm 0.5\%$ in ϕ_A and ± 8.5 ppb in δ_f. The propagation of the latter to the evaluation of volume fraction results in an estimated uncertainty of approximately $\pm 0.2\%$ in the volume fraction in the bubbly flow experiments, $\pm 3\%$ in the emulsion, and $\pm 3.5\%$ in the particle-laden flow. These values are the used to express the uncertainty of the measurements in Figs. 1.(d-f), 2.b, 3 and 4.

0.3.3 Effect of averaging

Lastly, we investigate the convergence of the averaging process to assess when additional averaging yields only marginal improvements to the result. In this work, averaging is done on spatial-domain images, rather than in K-space. To quantify the effect of averaging, 100 acquisitions of a pure water sample are collected after thermal steady-state is reached, yielding 100 pairs of amplitude and phase images. Progressive averages are then computed, resulting in $\mathcal{A}(x, y, N)$ and $\Theta(x, y, N)$, where the n-th frame corresponds to the average of the first n acquisitions. The progressive noise reduction $\mathcal{N}$ is then computed as the space-averaged standard deviation of the difference between the n-th and the n-th - 1 acquisition, normalized by the initial value (Eqn. 15). Results for both amplitude and phase images are reported in Fig. 8.

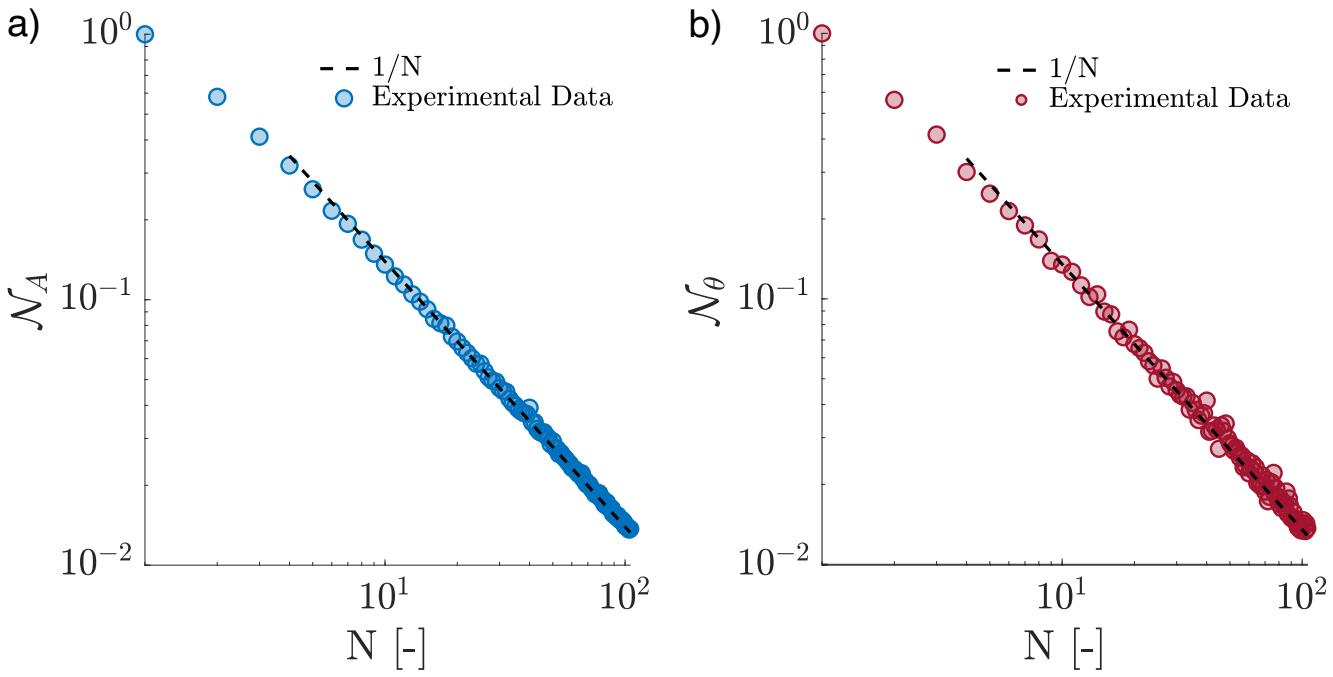

Figure 8: Convergence of the averaging process for both amplitude (**a**) and phase (**b**) images. In both cases, the normalized incremental contribution decreases following the expected trend 1/N.

$$\mathcal{N}_A = \frac{\overline{\sigma}(\mathcal{A}_n - \mathcal{A}_{n-1})}{\overline{\sigma}(\mathcal{A}_1)}, \qquad \mathcal{N}_\theta = \frac{\overline{\sigma}(\Theta_n - \Theta_{n-1})}{\overline{\sigma}(\Theta_1)}. \tag{15}$$

Since each new sample contributes to the running mean with weight $1/N$, the magnitude of the incremental update decreases as $1/N$. This metric therefore quantifies the regime in which further averaging yields only marginal improvements in the reconstructed images. Both amplitude and phase images scale with the same expected trend, thus allowing for a direct comparison of the two when possible.