

Chipiron

High-quality 1 mT MRI

Dimitri Labat*

Chipiron - 2/8 rue Gaston Rebuffat, 75019 Paris

(Dated: March 10, 2025)

Chipiron is developing a portable MRI device that is 10× cheaper than conventional systems, powered by an ultra-sensitive quantum detection system to address MRI accessibility. We introduce a new paradigm in medical imaging to detect deadly diseases earlier than ever and to closely monitor therapy effectiveness at scale. Everyone deserves easy access to MRI. In the future, undergoing an MRI will be as routine as a blood test.

HOW DO WE UNLOCK ACCESS TO MRI AT SCALE?

Context

Magnetic resonance imaging (MRI) is the most advanced medical imaging modality. It provides a wide variety of soft tissue contrasts at millimetric resolution in a few minutes of acquisition. Depending on the contrast (or sequence) used, one can visualize anatomic features, vessels, tumors, lesions, or even subtle differences in temperature or chemical composition in muscles and internal organs. Some MRI sequences even allow monitoring of metabolic or functional activity. For all these reasons, MRI is an essential tool for diagnosing many life-threatening diseases, for treatment follow-up (e.g., chemotherapy or neurodegenerative disease therapies), and for image-guided therapy and surgery.

Yet MRI remains largely inaccessible to 90% of the world’s population. MRI machines typically weigh about 5 tons, require a magnetically shielded room, expensive and complex maintenance, and trained staff to operate. They produce very strong magnetic fields using heavy, precisely engineered electromagnets. Higher magnetic field strength yields higher image quality in a given scan time for two main reasons:

- Protons inside the body align with the applied magnetic field, much like tiny bar magnets. The stronger the field, the more protons align, and the more signal is generated.
- The signal is detected with an inductive antenna, and an antenna’s sensitivity increases at higher frequency. Because MRI frequency is proportional to magnetic field ($\omega_0 = \gamma B_0$), a higher field produces a higher-frequency signal to which inductive antennas are more sensitive.

For decades, MRI manufacturers have pushed toward higher field strengths, as this is the primary driver of image quality. Most clinical MRI scanners operate at 1.5 T or 3 T (with research systems up to 7 T, fewer than 150 units worldwide).

However, pushing to ever-higher fields is what makes MRI machines so large, expensive, and ultimately not widely accessible. Therefore, the only way to democratize MRI is to achieve clinically relevant images at lower magnetic fields. This idea is as old as clinical MRI itself: since the early days, engineers have tried to build clinical-grade low-field MRIs. To date, no attempt has achieved routine clinical use at scale, for many complex reasons.

First (and obviously), it is very challenging to perform MRI at lower fields, for the reasons stated above: there is less signal and the detection is less sensitive. Consequently, the signal-to-noise ratio (SNR) is low. SNR is the key determinant of image quality: a high SNR yields a high-quality image with fine resolution in a short time. To address the lack of SNR at low fields, many strategies have been implemented. One can use a more sensitive, lower-noise detection system to improve SNR. One can also apply denoising software techniques such as EDITER [1] or deep-learning-based methods [2] to remove noise from the final images.

However, magnetic field strength is by no means the only driver of image quality. Images produced on 1.5 T scanners have improved tremendously since the 1990s, primarily thanks to advances in software and image reconstruction. The bottom line is that we need to (1) converge on the best possible hardware to perform MRI at low fields, and (2) make the best use of AI techniques for image enhancement such as super-resolution, advanced reconstruction, and denoising algorithms.

Low-field MRI has suffered for decades from a reputation of poor image quality. Yet what truly matters for clinical adoption is a combination of *perceived* image quality—which can be maximized through clever acquisition procedures—and sufficiently high *information* quality, which comes primarily from the MRI instrumentation and active noise cancellation techniques. There is often a misalignment between what clinicians perceive as high image quality and actual diagnostic power [3]. Low-field MRI has already demonstrated the baseline sensitivity and specificity for certain use cases. Some of the most promising use cases include:


- Population-wide screening for common deadly can-

cers such as breast or prostate cancer,

- Use in emergency settings for trauma and acute stroke,
- Use in intensive care units when patients cannot be moved (need for a portable MRI machine),
- Routine follow-up of patients to assess treatment effectiveness, particularly in Alzheimer’s disease (with new drugs entering the market) and oncology (chemotherapy monitoring).

At Chipiron, we believe that in 20 years low-field MRI will be as common as routine blood tests. It will be used for all the aforementioned applications and many more. Most importantly, it will unlock unprecedented capabilities in health monitoring and personalized medicine. Yet the hard question isn’t what low-field MRI *can* do, but rather which first use case will drive mass adoption. A good first use case is one where:

1. Low-field MRI has the required sensitivity and specificity — so it meets FDA thresholds and clinicians are comfortable using it. It goes without saying that 10 mT MRI will never replace 3 T MRI; we will still need 3 T scanners whenever they are available, especially for the most challenging cases. Low-field MRI will be used where it provides sufficient information to make an informed clinical decision.
2. High-field MRI cannot be used — whether due to limited availability, patient/environment incompatibility, or cost. One never wants to be directly compared to high-field MRI because users would always prefer the higher-field option if available.
3. It is clearly better than the current alternative. For example, consider prostate cancer screening. In theory it’s a good first use case, because compared to the gold standard (PSA testing plus biopsy) it is non-invasive and potentially cheaper. However, despite promising early results, low-field MRI is clearly not more sensitive — at best on par with the gold standard. Even high-field MRI is not definitively better than biopsy. For this reason, despite its benefits, most clinicians would be reluctant to use low-field MRI for fear of missing a potentially deadly lesion.
4. It offers a rapid return on investment and a reliable revenue source for the institution. This is tricky for many diagnostic applications. The solution must be an order-of-magnitude cheaper *and* more effective than the current alternative to justify the investment. It is sometimes tempting to assume that because a solution is somewhat effective, cheaper,



and less invasive/more comfortable for the patient, clinicians will adopt it — but this assumption has often proven wrong.

We are fortunate to have an established competitor already on the market: Hyperfine. Their first product, the Swoop, was FDA-cleared in 2021, followed by a \$580 M IPO at the end of that year. It was initially marketed for intensive care units and acute stroke diagnosis [4]. Since then, their image quality has improved dramatically, mainly thanks to better AI reconstruction and denoising, as illustrated in Fig. .

Despite these impressive improvements in image quality, a true product–market fit for the Swoop device has yet to be found. According to the latest Q3 2024 financials [5], Hyperfine reported just over \$3 M in sales for that quarter, corresponding to roughly 10 installed units. The company recently pursued another use case: monitoring patients with Alzheimer’s disease (AD), specifically looking for ARIA imaging abnormalities related to new AD treatments targeting amyloid plaques. Initial results [6] suggest that the Swoop could be as effective as high-field MRI for tracking white matter hyperintensity volume over time. This new direction is also interesting from a business perspective, as the imaging is performed in the context of therapy rather than large-scale screening, potentially leading to more consistent revenue. Hyperfine is doing very important work pioneering new low-field practices in a particularly challenging and conservative market.

How will Chipiron win

There are clear reasons why low-field MRI has not yet achieved routine clinical use. Many of these reasons still stem from the low field strength itself. Operating at 64 mT, for instance, means using a permanent magnet that requires no electrical power or cooling. However, such magnets are heavy (over 600 kg for Hyperfine’s system) and very inhomogeneous (limiting the usable sequences), and as a result they are confined to extremity imaging in a tight geometry. Counter-intuitively, we believe that for low-field MRI to truly break into clinical practice, we

need to go to even lower fields. Ultra-low fields (below 10 mT) enable fundamentally new approaches to MRI for several reasons:

- Very low B_0 fields can be generated with a simple copper electromagnet, replacing the heavy permanent magnet. The result is a very open geometry that can accommodate full-body imaging of patients of any size, with easy access to the patient. This is impossible with permanent magnets or traditional high-field systems.
- Lower fields also mean weaker eddy currents and susceptibility artifacts. This makes the system compatible with most other instruments — for instance, surgical tools or other imaging modalities.
- Tissue contrast can actually increase at lower fields, allowing better discrimination between tissues. This has been shown in prostate cancer screening [7] and acute stroke diagnosis, at fields from a few hundred μT to a few mT.

Of course, MRI at 64 mT was already a challenge; at 10 mT it will be even greater. But we believe it is the only path to making MRI truly ubiquitous, and we are determined to make it happen — with the help of an ultra-sensitive quantum detection system: the superconducting quantum interference device (SQUID).

TECHNOLOGY

Ultra-low field MRI

Ultra-low field MRI (ULF MRI) simply refers to MRI performed at very low magnetic fields. The typical B_0 field we work with is below 10 mT, about a thousand times lower than a standard clinical field (taking 3 T as a reference). Before diving into the subtleties of very low fields, let's briefly review the principle of MRI and why it is such an extraordinarily powerful imaging technique.

MRI is based on the phenomenon of nuclear magnetic resonance (NMR). It leverages the spin properties of nuclei that make up molecules in the body, in particular hydrogen. All hydrogen nuclei can be thought of as tiny magnets, and we apply a sequence of magnetic fields to them. By probing their response to these fields, we deduce local tissue properties and construct a 3D image of the soft tissues inside the body. A typical MRI procedure can be broken down into the following steps:

- The body is placed in a static magnetic field B_0 . The spins of hydrogen nuclei are distributed into two quantum energy states separated by energy $\Delta E = \hbar\omega_0$, with $\omega_0 = \gamma B_0$. This creates a net polarization P of the sample aligned with B_0 , proportional to B_0 .

- We apply radiofrequency (RF) pulses at frequency ω_0 to the body. This perturbs the spin system out of equilibrium: the polarization P , initially aligned with B_0 , is tipped into the B_0 transverse plane.
- The spin system returns to equilibrium with characteristic relaxation times T_1 and T_2 , emitting RF signals at ω_0 in the process. During this relaxation phase, we also apply spatially varying *gradient* magnetic fields to encode spatial information in the signal. We detect these signals with a receive coil and reconstruct maps of the local T_1 and T_2 relaxation times, which are characteristic of the tissue. These maps form the basis of MRI images.
- The process is repeated N times — exciting the sample, applying gradient fields, and measuring with the receive coil. By averaging the results, the signal-to-noise ratio improves by a factor \sqrt{N} .

If you are to remember only one equation, it should be:

$$\omega_0 = \gamma B_0, \quad (1)$$

which tells us that the central MRI frequency ω_0 [8] is proportional to the magnetic field B_0 via the gyromagnetic ratio γ of the proton.

Because $\gamma/2\pi \simeq 40 \text{ MHz/T}$, we have $\omega_0 \approx 120 \text{ MHz}$ at 3 T but only $\approx 40 \text{ kHz}$ at 1 mT. Thus, when operating at fields between 1 and 10 mT, we deal with frequencies on the order of 100 kHz. This $\sim 100 \text{ kHz}$ frequency range is interesting for a couple of reasons:

- 100 kHz is extremely slow compared to modern RF electronics (which operate in the GHz range). This makes it easy to design feedback and readout electronics to suppress noise and ensure system stability.
- At such low frequencies, common MRI artifacts caused by metal (susceptibility artifacts and eddy currents) are greatly reduced.
- *Body noise*, which is the dominant noise source at high fields, becomes lower than the detector's noise at ultra-low fields. This means there is room to improve SNR by lowering the intrinsic noise of the detection system.

Detecting tiny NMR signals

The fundamental question for ultra-low field MRI is how to make the most of the very small amount of available signal. As mentioned in the introduction, two independent factors make your life harder at ultra-low fields.

Firstly, and most importantly, higher field means more available signal. Understanding why requires a bit of

quantum physics (details are given in the appendix, but here is the gist). Hydrogen nuclei (protons) have spin $1/2$ and as a consequence a magnetic moment. When the body is placed in a magnetic field B_0 , these protons can occupy one of two energy states with energies $E_{1,2} = \pm E_0/2$, where $E_0 = \hbar\omega_0$ and $\omega_0 = \gamma B_0$. If the temperature were zero, we would expect all protons to occupy the lowest energy state $E_1 = -E_0/2$. At room temperature $T \sim 300$ K, however, the protons are distributed between the two states. Because the thermal energy $k_B T$ is much larger than $\hbar\omega_0$, the two states are almost equally populated, with just a slight excess in the lower energy state. This yields a tiny net polarization in the direction of B_0 . The polarization of the sample can be approximated as

$$P = \frac{N(E_2) - N(E_1)}{N(E_2) + N(E_1)} \simeq \frac{\gamma \hbar B_0}{2k_B T}. \quad (2)$$

Aside from increasing the magnetic field, nothing in this formula can be tuned (the subject remains at body temperature T). At $B_0 = 3$ T, we get $P \simeq 10^{-5}$ — an incredibly small fraction (only about 10 protons in a million contribute to the NMR signal). At $B_0 = 10$ mT it is even worse: only ~ 30 protons per billion contribute. This weak polarization is the fundamental curse of NMR and explains why it is so challenging.

But our problems do not end there. To detect the magnetic signal $B_2(t)$ emitted by the sample, the usual approach is to use an inductive antenna — a copper loop of area S — in which the NMR signal induces a current by Faraday’s law:

$$Ri = -\frac{d\Phi}{dt}, \quad (3)$$

where i is the current induced in the loop, R is the loop’s resistance, and $\Phi = B_2 S$ is the magnetic flux of the signal through the loop [9]. If we consider a signal of frequency ω_0 (tuned such that $\omega_0 = \gamma B_0$), the relation between the induced current and the polarization field B_0 is:

$$i = \frac{1}{R} \gamma B_0 B_2 S. \quad (4)$$

A sensitive antenna is one that produces a large current i for a given field B_2 . This equation shows that inductive antennas are much more efficient at high B_0 . If we lower the field from 3 T to 10 mT, the induced current is 300 times smaller. In summary, compared to high field MRI, we have $\sim 300\times$ less signal and the conventional detection method is $\sim 300\times$ less sensitive, amounting to a problem of size 10,000 to solve.

There is no way to overcome the low signal itself without raising the field or cooling the patient (which is impractical). Instead, we must work on the detection system — and this is where superconducting quantum interference devices (SQUIDs) come into play. SQUIDs are

extremely sensitive magnetometers with a flat frequency response from DC up to a few hundred MHz. The core of a SQUID is a micrometer-scale superconducting ring (in our case made out of Niobium) interrupted by two Josephson junctions. We provide a detailed explanation of superconductors and SQUID operation in the supplementary material. SQUIDs belong to the broad family of ”quantum sensors” [10] that includes SQUIDs, atomic vapor magnetometers (of the SERF or scalar type), diamond NV center magnetometers, kinetic inductance detectors, etc.

Because SQUIDs are so small, they are rarely used alone to measure far-field signals. Usually, they are integrated into a larger detection system in a ”current-sensing” mode where the SQUID is coupled to a macroscopic antenna called a flux concentrator, similar to the receive coil in a regular MRI. The key difference is that a SQUID antenna is usually *gradiometric*, meaning it comprises two coils wound in opposite directions to reject distant uniform noise. Our approach at Chipiron, compared to older SQUID MRI setups, is the invention of a new type of SQUID-based volume gradiometer. With this geometry, SNR improves by a factor of 3 to 10, unlocking the possibility of performing ultra-low field MRI outside of a magnetically shielded room (details are provided in the appendix).

In current-sensing mode, a SQUID can be thought of as a very high-gain, low-noise current amplifier for the MRI antenna. The critical difference from a conventional MRI receive coil is that instead of using a low-noise voltage amplifier to detect the coil’s voltage, we use the SQUID as a current-to-voltage transducer to amplify the coil’s current [11]. This subtle point is the core reason why a SQUID is the best amplifier at ultra-low fields. Indeed, if you use a cooled pickup coil with sufficiently low resistance R relative to its inductive reactance $L\omega$ (such that $R \ll \omega L$), the gain of the SQUID detection chain becomes frequency-independent. This is in contrast to a regular MRI antenna, where you amplify voltage and inevitably lose sensitivity as frequency drops, no matter what the coil’s resistance is.

AI image enhancement

Ultra-low field MRI will never replace conventional high-field MRI [12]. In fact, it unlocks access to MRI in situations where high-field systems are not feasible; in this sense, low field is an extension of high field.

As mentioned earlier, the only way to enable mass adoption of low-field MRI is to make clinicians comfortable with its inherently lower image quality. With far less signal available, one must make compromises. There are four parameters that govern image quality, all of which depend on the global SNR of the image:




FIG. 1. Schematic of the different parts of the SQUID detection system. The image depicts a room-temperature volume gradiometer coil that functions as a flux concentrator, with compensation coils on either side to minimize far-field noise. The coil is coupled to the SQUID sensor via a shielded cable at 4 K to capture and amplify the detected signal.

- Acquisition time t ,
- Contrast-to-noise ratio (CNR),
- Spatial resolution a ,
- Signal-to-noise ratio per voxel (SNR_v).

Acquisition time can be adjusted by changing the number of repetitions N of the MRI sequence. By repeating the sequence multiple times, random noise averages out and SNR increases. However, it is crucial to keep acquisition times within clinically acceptable ranges so as not to disrupt workflow. To that end, we keep t below 5 minutes per sequence. The CNR depends on field strength (like SNR, but more subtly). It is expected that at very low fields, some contrasts improve, as shown in the literature for acute stroke (in vivo) [13] and prostate cancer imaging (ex vivo) [14]. But once the field is chosen, CNR is fixed and not tunable.

This leaves two remaining parameters: spatial resolution and SNR per voxel. While most MRI scans easily achieve sub-millimeter resolution, we target a more modest resolution of $2 \times 2 \times 4 \text{ mm}^3$. According to feedback from a panel of 100+ radiologists we surveyed, this is sufficient for the initial use cases we envision. With acquisition time and resolution fixed, the SNR per voxel is also fixed. Through simulations and phantom scans, we estimate we will achieve an SNR_v between 30 and 50

per voxel — significantly lower than at high fields, but sufficient for the first use cases.

The important question now is: are clinicians comfortable working with this baseline image quality? This is where image enhancement techniques come into play [15]. Since the FDA clearance of Hyperfine’s first Swoop device in 2021 and the advent of low-field MRI in clinical practice, there has been a surge of interest in techniques to transform low-field images into high-field-like synthetic images. These techniques are especially effective for estimating quantitative parameters. For example, super-resolution algorithms have been applied to 64 mT images to quantify white matter hyperintensities in the progression of Alzheimer’s disease [16]. Qualitative AI enhancements are another story. A recent study [17] — which caused a stir in the low-field community — introduced a diffusion model to synthesize 3 T-like images from a 55 mT system capable of brain and body imaging. While the perceived image quality is flattering, it remains unclear if these images have clinical value. The general problem with any AI-based image enhancement is twofold: (1) synthesize a clinically accurate image (avoid hallucinations and preserve true anatomy), and (2) remain faithful to the patient’s specific anatomy. The second point is particularly tricky, as AI methods tend to produce an “average” human anatomy, which could obscure small but important patient-specific details or lesions. A good way to mitigate this is to incorporate prior information

about the patient when available (e.g., medical history or recent high-field MRI). In many cases, such prior data exist; for example, in intraoperative low-field MRI for neurosurgery, a pre-operative high-field MRI provides a very recent image of the patient’s brain. Having a series of earlier scans also allows acceleration of follow-up scans using compressed sensing techniques [18].

We have recently begun developing a pipeline to generate synthetic 3 T-like images from ultra-low field acquisitions. Before we obtain low-field clinical data, we are creating an artificial dataset by degrading high-field images to the anticipated SNR and resolution [19]. Other techniques for synthetic MRI image generation include generating fully simulated data [20][21] or using a combination of scans with different contrasts, orientation and resolution [22]. Using a sample from the M4Raw 0.3 T dataset [23] as a starting point, we downsample (crop k-space) to adjust resolution and add white Gaussian noise to achieve the desired SNR. This simulated dataset is only a useful starting point and not an entirely accurate representation of actual low-field data, partly due to the strong field dependence of contrast [24]. Our initial efforts focus on super-resolution techniques that require no prior information. We have started benchmarking state-of-the-art approaches — such as the SRDenseNet method described by de Leeuw den Bouter *et al.* [25] by adapting it to 3D — and a 3D V-Net architecture (a volumetric U-Net variant) with a sliding window for added context in training and inference. On the reconstruction side, we are interested in methods robust to variability in the exact sampling pattern [26] to ease integration into the clinic.

PRODUCT

Prototypes

We started from an empty lab space three and a half years ago. Since then, we have built two ultra-low field MRI prototypes operating from 1 to 10 mT. Our first 1 mT prototype, “Pacific,” is shown in Fig. .

The Pacific prototype allowed us to acquire our first images and overcome many hardware hurdles (some are explained in detail in the appendix): taming noise from switching gradient amplifiers, stabilizing the SQUID readout chain, testing multiple acquisition consoles, and generally optimizing the system. We then built a second MRI system, *Atlantic* (Fig.), which serves two purposes:

- Integrate a cooled pickup coil into the detection chain. The Pacific prototype uses a room-temperature pickup coil. Cooling the coil to ~ 50 K can, according to simulations, reduce noise ten-fold and dramatically boost SNR. We built a custom cryostat for the pickup coil, which is currently




FIG. 2. The Pacific 1 mT prototype. The goal of this first device is to demonstrate that SQUID-detected 1 mT MRI works on non-clinical phantoms (like fruits and ionic-solution phantoms) and to assess the best image quality we can achieve at this field. The B_0 field is generated by a 40 A current through the four vertical coils of the cage-like structure (a). The sample is placed between printed circuit board gradient coils (b) driven by AE Techtron linear amplifiers to generate up to 5 mT/m gradients with less than $100 \mu\text{s}$ rise time. The MRI signal is picked up by a room-temperature copper coil around the sample and sent inside the cryostat (c), where it is amplified by a low- T_c SQUID (SQ2600, STAR Cryoelectronics). The signal is then read out by room-temperature electronics in the rack (d).

under test to measure the SNR improvement.

- Test different B_0 field strengths up to 10 mT. Lower fields yield higher T_1 contrast but lower SNR. There is an optimal field that may depend on the tissue. Since we use a resistive magnet, we have the flexibility to test different fields and settle on the optimal field for the investigational device.

Q4 2024 marked the end of the prototyping phase at Chipiron: we gathered all data needed to finalize the specifications of our first clinical device.

First images

All acquisitions to date have been performed with a room-temperature pickup coil coupled to a 4 K SQUID detector (SQ2600, STAR Cryoelectronics), and a Pure Devices spectrometer for data acquisition. The transmit coil is a tuned saddle coil, and the receive coil is a tuned gradiometric solenoid. The B_0 field is 1.0 mT for the Pacific prototype and 10.0 mT for the Atlantic prototype, stabilized by a PID feedback loop on the field measurement.

From Fig. 10, we can estimate the SNR improvement needed to reach clinical image quality. We already have the required acquisition time and spatial resolution, so the only remaining improvement is:



FIG. 3. The Atlantic 10 mT prototype. Like Pacific, the B_0 field is generated by a 40 A current, here through a smaller copper solenoid (a). The sample is placed between the gradient plates and main coil (b). The MRI signal is picked up by a room-temperature copper coil around the sample and sent inside the cryostat (c) where it is amplified by a low- T_c SQUID (SQ2600, STAR Cryoelectronics). In this case, the cryostat is open, revealing the gold-plated structure that cools the SQUID. In later iterations, this room-temperature pickup coil will be cooled with the custom cryostat described in Fig. . The signal is then read out by room-temperature electronics in the rack (d).




FIG. 4. Our first cryostat housing a cooled pickup coil. It comprises a metallic 4 K stage (on the right) hosting the SQUID sensor, and a vacuum-sealed G10 enclosure at 50 K housing the pickup coil (beige sleeve on the left). This cryostat can accommodate an 80 mm diameter room-temperature sample for some initial phantom tests. We expect to produce our first MRI images with this system by the end of Q1 2025.




FIG. 5. A sketch of the Chipiron investigational device (ID) currently being assembled. This design minimizes technological risk for acquiring the first clinical images, and does not include all features planned for the first product. In particular, the solenoid geometry shown here will not be kept for the product, as we prefer a more open Helmholtz-like design. This ID is dedicated to head imaging only, whereas future product iterations will accommodate full-body imaging.



FIG. 6. Design of a later Chipiron product iteration that is portable and intended for extremity imaging.

- The global signal-to-noise ratio, currently around 5–10. Ideally we want an SNR of 50, so conservatively we need a 10-fold increase.
- The phantom used is a 5 cm sphere. Taking 2000 cm^3 as an upper bound for human cranial volume, we need about a 30-fold SNR increase to cover that field of view.

To image a human brain in 5 minutes at $2 \times 2 \times 4 \text{ mm}^3$ resolution, we thus need roughly a $300\times$ SNR increase. We plan to achieve this factor of 300 through:

- Cooling the detection antenna, for a $\sim 10\times$ SNR gain.

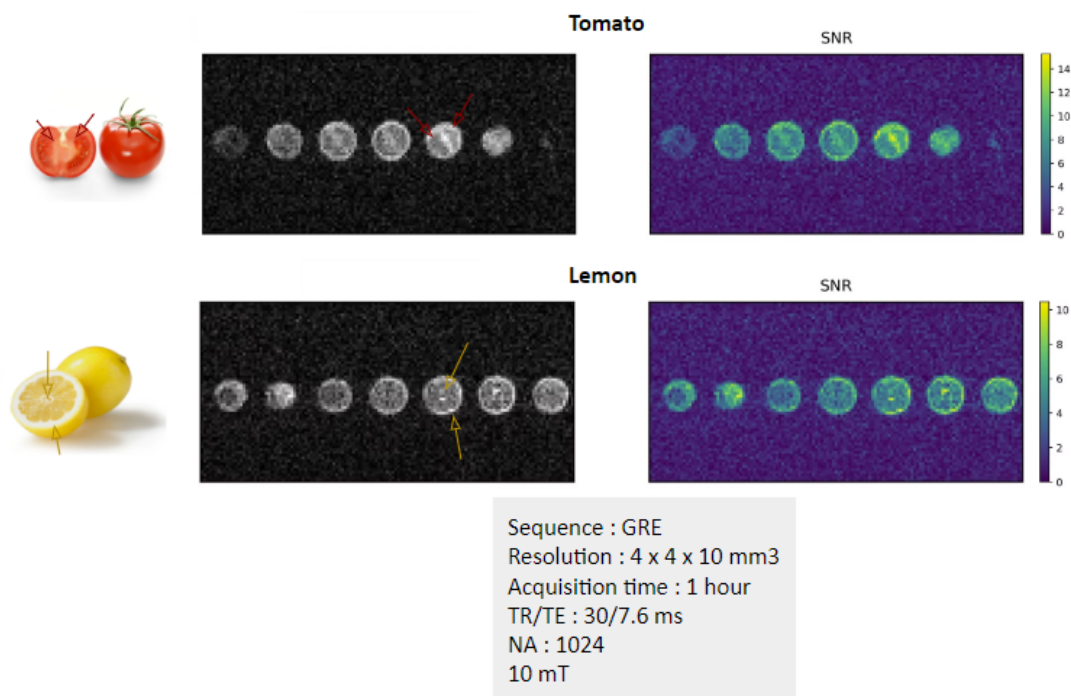


FIG. 7. GRE images of a lemon and a tomato acquired at 10 mT with our SQUID detection system. Sequence parameters: number of averages = 1024 for a 1 hour acquisition; spatial resolution $\sim 4 \times 10 \times 10 \text{ mm}^3$. These first biological images show internal structures of the fruits.




FIG. 8. GRE images of a garlic clove acquired at 10 mT with our SQUID detection system. Sequence parameters: number of averages = 154 for a 1 hour acquisition, and 1848 for a 12 hour acquisition, respectively. The spatial resolution is $\sim 2.5 \times 2.5 \times 10 \text{ mm}^3$. The images clearly show how longer acquisition time improves SNR.

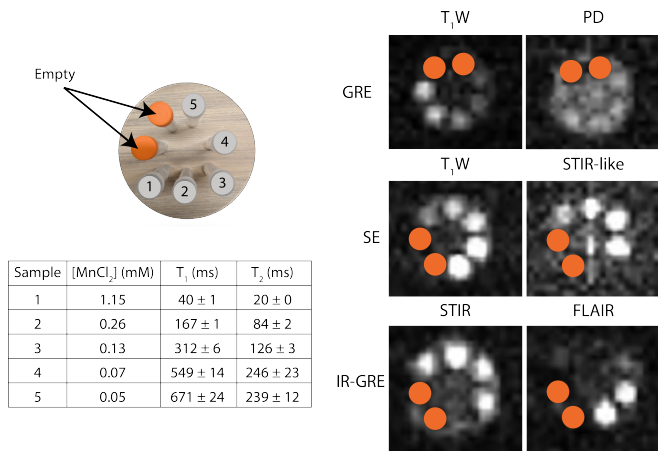


FIG. 9. GRE images of a contrast phantom acquired at 10 mT with our SQUID detection system, giving an early indication of how contrasts may appear at lower fields.

- Implementing high-coupling impedance matching cryo-transformers, for another $\sim 10\times$ SNR gain. These are under testing this quarter.
- Accelerating acquisition via a combination of Gaussian undersampling, echo planar imaging (EPI), and partial Fourier, for a $\sim 5\times$ effective SNR gain. We have demonstrated each component individually; they will be fully integrated into the final investigational device by summer 2025.

These improvements multiply to an SNR increase of $\sim 500\times$, overshooting the target of 300. The estimates used are conservative; moreover, we have not accounted for active noise cancellation (ANC), which early results suggest could provide an additional 5–10 \times SNR boost. We are confident that we will achieve the baseline image quality needed for our first clinical images by early 2026. This baseline image quality will be the foundation for the AI image enhancement pipeline we have begun developing to generate 3 T-like images from our 10 mT device.

Go-to-market elements

Chipiron’s primary goal is to install a first investigational device in a clinical setting by the end of 2025. Until that point, all our scans have been of phantoms, conducted in our lab by Chipiron staff. It will be the first time our device is used on actual patients by MRI technologists and radiologists in a clinical environment. The purpose of this investigational device (which is not yet a commercial product) is to evaluate the quality of our first brain images. To that end, we have set up collaborations with leading clinical institutions in France to deploy three devices in parallel. In our protocol, a selected cohort of

patients coming for standard high-field MRI will also undergo a Chipiron ultra-low field exam. This will allow us to build a database of paired low-field/high-field images to (1) train a dedicated AI pipeline to boost the image quality of the mT images toward high-field quality, and (2) assess the sensitivity and specificity of our sequences for each use case. With this information, we will be able to choose a first clinical application and submit an FDA 510(k) by the end of 2027 to enter the US market.

ACKNOWLEDGEMENTS

The work presented in this white paper was conducted by Chipiron’s R&D and product teams, which as of March 2025 consist of: Nour Azzoug, Zineb Belkacemi, Yacine Belkhodja, Pierre Chavanne, Lola Courty, Romain Couvreur, Bastien Dassonneville, Andy Decool, Pierre Emerich, Ricardo Ferreira, Guillaume Daval-Frerot, Marco Fiorito, Nina Gidel-Dissler, Eva Grimaldi, Sahil Jagga, Alexandre Jaoui, Sucharitha Kommu, Dimitri Labat, Ijee Mohanty, Abel Rangel Trejo, Mathis Rioufol, Isabelle Saniour, and Mustafa Utkur.




FIG. 10. A first image of a NEMA phantom acquired in 3.7 minutes (target acquisition time <5 min) at 10 mT with target resolution $2 \times 2 \times 4 \text{ mm}^3$. The last remaining step is to improve SNR and field of view for in vivo human head imaging.

SUPPLEMENTARY MATERIAL

Signal intensity in MRI

The conventional paradigm in MRI is to use higher and higher field strengths, for two main reasons. The first is that a higher field means a higher sample polarization. This polarization represents the fraction of protons that contribute to the MRI signal. The higher the field, the higher the polarization, and the more signal is emitted by the protons — leading to better images and shorter scan times.

To understand this, consider a single proton in a static field B_0 . Protons are spin $1/2$, so the Hamiltonian of the system is:

$$\hat{H} = -\hat{\mu} \cdot \vec{B} = -\mu_0 B_0 \hat{\sigma}_z, \quad (5)$$

where μ_0 is the proton's magnetic moment, B_0 is the static field, and $\hat{\sigma}_z$ is the Pauli Z matrix:

$$\hat{\sigma}_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (6)$$

From this, we define the Larmor frequency ω_0 such that $-\mu_0 B_0 = \hbar\omega_0/2$. We also define the proton gyromagnetic ratio $\gamma = 2\mu_0/\hbar$, so that $\omega_0 = \gamma B_0$. We are left with a simple two-level system with an energy splitting $\Delta E = \hbar\omega_0$ between the two states. Now, instead of a single proton, we have a macroscopic collection of them that we need to distribute between these two energy states. If the temperature were very low, these spin- $1/2$ particles would follow a Fermi-Dirac statistic. But of course the subject is at room temperature, so we just follow a very simple Boltzmann distribution:

$$P = \frac{N(E_2) - N(E_1)}{N(E_2) + N(E_1)} \quad (7)$$

$$\simeq 1 - \exp\left(-\frac{\gamma\hbar B_0}{2k_B T}\right) \quad (8)$$

$$\simeq \frac{\gamma\hbar B_0}{2k_B T}. \quad (9)$$

At a field of $B_0 = 3$ T, we get $P \simeq 10^{-5}$ — which is incredibly small: effectively only about 10 protons in 1 million are contributing to the NMR signal! At a field of 10 mT it is even worse, only ~ 30 protons per billion contribute to the signal. All of this is because the two spin states are almost equally populated. This fundamental limitation is the great curse of NMR and explains its challenges. Usually in MRI you have to repeat the sequence N times to average out noise and increase SNR. In high-field experiments where the signal is larger, you can get away with $N = 1$ or just a few repetitions, whereas in low field N may need to be in the hundreds.

To understand the influence of various parameters on SNR, it is useful to define two quantities: the global SNR

of the image, and the SNR per voxel, or SNR_v . Considering an acquisition with isotropic spatial resolution a (voxel volume a^3), number of repetitions N , total acquisition time t , and field-of-view length L (volume L^3), the scaling laws are:

$$\text{SNR} \sim \sqrt{N}, \quad (10)$$

$$\text{SNR}_v \sim a^{-3}, \quad (11)$$

$$t \sim N \sim L^3. \quad (12)$$

From these relations we can draw several important conclusions:

- SNR per voxel scales with the inverse cube of the resolution — a very steep dependence. In practice, this means that going from 2 mm to 1 mm resolution doesn't cost a factor of 2 in SNR, but a factor of 8. This is why we start with a conservative 2 mm resolution initially, and only decrease it gradually as SNR improves in future product iterations.
- Repeating the acquisition N times can boost SNR by a factor \sqrt{N} . This is useful but cannot be pushed too far. For example, doubling the scan time (N) yields about a 40% SNR increase, which might be acceptable; however, a $10\times$ longer scan only triples SNR, which is rarely worth it. Conversely, if you can double SNR per repetition through hardware or reconstruction improvements, you can cut the number of repetitions (and scan time) by a factor of 4 — a huge speedup. This is why we focus heavily on improving SNR through hardware, denoising, and reconstruction.
- Not surprisingly, acquisition time scales with volume of coverage. This is one reason we must work hard to go from imaging a small fruit to imaging a human head!

Gradients and images

So far, we have only explained how protons emit an NMR signal across the whole sample. To form an image, we need to know where within the sample the detected signal is coming from. This is where gradient coils come into play. On top of the homogeneous B_0 field, we superimpose small magnetic field variations $\delta B(x, y, z)$ that vary linearly along the three spatial directions. In a commercial 1.5 T MRI, the gradient strength is on the order of 100 mT/m. The gradient field is produced by three sets of *gradient coils* installed near the main B_0 coil.

Consider 2D imaging, where the 3D image is composed of a stack of 2D slices. The gradient coils produce a small field $\delta B(x, y) = G_x x + G_y y$ on top of the homogeneous B_0 . The field experienced by a proton at position (x, y) in a slice is $B_0 + \delta B(x, y)$. Following our earlier discussion,

this proton will precess at a frequency $\omega(x, y) = \gamma[B_0 + \delta B(x, y)]$ that depends on its position (x, y) .

To construct an image, we need to determine from which *voxel* (3D pixel) the signal originates. That is precisely what the gradients enable. Typically, during the sequence, one first applies a short *phase-encoding* gradient G_y . After this, protons at different y positions accumulate different phase offsets (because they spun at different frequencies during the phase encode). Then a *frequency-encoding* gradient G_x is applied while the NMR signal is acquired. This procedure is repeated N_p times with varying G_y strengths, where N_p is the number of phase-encoding steps. In the end, the phase of the NMR signal encodes the y position of the voxel, while the signal frequency encodes the x position. A 3D image is obtained by repeating this for N_s slices, using the G_z gradient and a selective RF excitation pulse to target each z slice.

The duration, strength, and order of the gradient and RF pulses form what is called an *MRI sequence*. Different sequences yield different image contrasts. In fact, certain sequences will highlight tissues with long T_1 (we call these T_1 -weighted images). In a CT scan, contrast is determined by a single parameter (X-ray absorption), but MRI contrast can depend on many parameters whose influence is tuned by the sequence: T_1 , T_2 , proton density, magnetic susceptibility, water diffusion tensor, etc. Recently, there has been growing interest in *quantitative* sequences like magnetic resonance fingerprinting that produce maps of local parameters (e.g., T_1 and T_2) instead of qualitative weighted images. The versatility of MRI is one of its greatest strengths and the main reason it holds such a unique place among medical imaging techniques.

Simulated clinical images

To estimate the SNR required to produce diagnostic-quality images, we performed realistic simulations of what clinical images would look like at various target SNRs.

The simulation is based on the 630×630 knee image shown in Fig. 13, with an approximate resolution of 0.3 mm.

The image is then processed as follows:

- Crop the top and bottom blank regions.
- Lower the resolution to the target by reducing the k-space extent.
- Apply a decay across the k-space lines corresponding to the normalized decay from real FID acquisitions. This step has minimal impact since the chosen readout window is short and placed near the start of the signal (start time 3 ms, readout window 2 ms).




FIG. 11. The same brain imaged with different MRI sequences. (a) T_1 -weighted. (b) T_2 -weighted. (c) FLAIR. (d) FLAIR with contrast enhancement. *Liu et al., Tsinghua Sci. Technol.* **19**(6): 578–595 (2014).

- Add Gaussian noise after several normalization steps:
 1. Scale to match the noise level observed after a long decay in a real FID (noise defined as whatever is measured after 1500 ms).
 2. Scale by the ratio of the target number of repetitions to the number of repetitions in the measured FID (accounting for total scan time).
 3. Scale by the ratio of the estimated number of protons in the target knee slice to that in a small reference sample.
 4. Scale by the ratio of the peak k-space value in the original image to that of the measured FID at $t = 0$.
 5. Divide by an arbitrary SNR boost factor given as input.
- Zero-fill k-space to artificially increase resolution.
- Apply Gaussian convolution to smooth zero-filling artifacts.

We simulated acquisitions for 4 scenarios corresponding to expected results over the next few development phases, shown in Fig. 14.

Active noise cancellation

(Author: Zineb Belkacemi.)

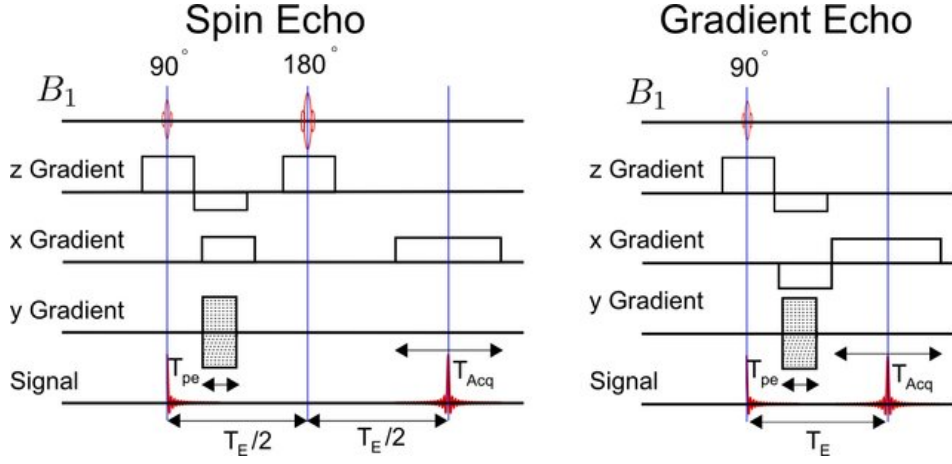


FIG. 12. Field intensity vs. time for spin-echo and gradient-echo sequences, the two main families of MRI sequences. B_1 is the RF excitation field. A spin-echo sequence uses an additional pulse to rephase the spins (hence the "echo", whereas in a gradient echo sequence, the echo is formed by successive dephasing and rephasing of the frequency-encoding gradient. Spin-echo sequences were one of the earliest MRI developments. Gradient echo sequences tend to be faster but are more sensitive to artifacts. (Adapted from the PhD thesis of Karl Edler.)



FIG. 13. Reference knee image used for our simulation.

Definitions Due to the low signal and desire for portability, our ultra-low field MRI system is highly susceptible to noise — both internal (e.g., Johnson-Nyquist noise) and external (especially in unshielded environments). Reducing external electromagnetic noise can be done at several points in an MR sequence, either in analog hardware (compensating magnetic field noise with coils, subtracting noise from the analog signal) or via numerical post-processing (denoising the MR signal with data-driven methods or AI). Here we summarize current and planned efforts to reduce EMI noise in post-processing. We focus on methods using external antennas [27][28][29][30]. The general idea behind these methods is to use external coils placed around the MRI (and its receive coil) to sample the ambient EMI noise during the sequence, then digitally remove that noise from the receive coil signal. Of course, one must first establish a




FIG. 14. Top to bottom, left to right: simulated knee images corresponding to anticipated improvements every 6 months. The expected acquisition times, cumulative SNR gains, and achieved resolutions are (24 h, $1\times, 4\times 4\times 8\text{ mm}^3$), (6 h, $4.5\times, 3\times 3\times 6\text{ mm}^3$), (1 h, $58.5\times, 2\times 2\times 4\text{ mm}^3$), and (10 min, $643.5\times, 2\times 2\times 4\text{ mm}^3$), respectively.

mapping between the noise in the external coils and the noise in the receive coil.

Formally, let S be the signal from the receive (R_x) coil, and N_1, \dots, N_c the noise signals acquired by c surrounding coils. Following the original EDITER paper [31], we assume the receive signal S is the sum of the "pure" MR signal S^* and the EMI noise N :

$$S = S^* + N. \quad (13)$$

Here, other noise sources are lumped into the "pure" MR signal S^* . To recover S^* , we want to find a relationship $N = f(N_1, \dots, N_c)$ that allows us to remove N from S .

Two categories of methods exist for external coil-based

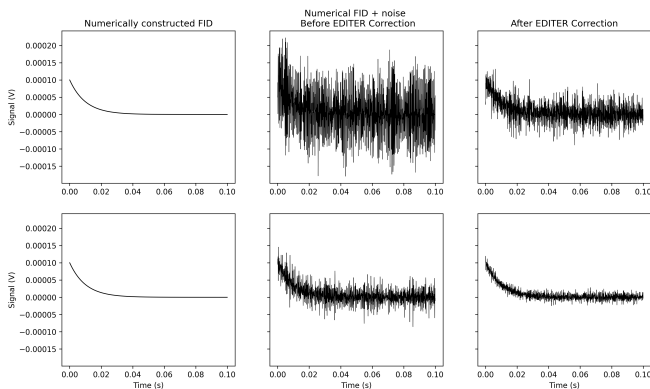


FIG. 15. Performance of the EDITER denoising algorithm on synthetic signals after demodulating noise signals. Top: a single-shot numerical FID ($T_2 = 0.01$ s, amplitude 0.1 mV). Bottom: averaged over 10 shots. (In both, red = original signal, purple = EDITER-corrected signal.)

noise cancellation:

- The first category requires calibration data acquired outside the receive window ($S^* = 0$). This can make it easier to approximate f explicitly for noise removal during imaging. However, it requires a “dead time” to acquire calibration data either before the MRI sequence or between repetition times, thereby increasing the minimum T_R .
- The second category does not require calibration data and can be applied on the fly to the MR signal, in theory making it compatible with any sequence. These methods often assume low correlation between S^* and N [32] or isotropic noise [33].

Preliminary results First, we evaluated external coil noise cancellation outside of the MRI setup.

We generated a dataset of noise signals using 7 coils: one central coil as the receive (R_x) coil, and $c = 6$ surrounding EMI coils. We tested two methods:

- A linear convolution-based method (EDITER), which is calibration-free. EDITER was tested by artificially adding a numerical FID to the R_x coil signal.
- A deep learning convolutional neural network (CNN) which requires calibration data for training. The network was trained on part of the simulated dataset and tested on a separate part.

So far, both methods achieve similar results in terms of residual noise amplitude, though EDITER’s success appears more dependent on the initial SNR. Further tests with poorly synchronized and/or differently processed signals (e.g., different tuning/matching, demodulation, filtering, etc.) are ongoing.

One way to quantify the performance of noise cancellation algorithms such as EDITER is to compute the ratio of remaining noise in the corrected signal to the original noise for a given added numerical FID. More precisely, letting N be the noise at the receive coil and adding an FID S^* as the pure MR signal, we compute:

$$R = \frac{\text{std}(S_{\text{EDITER}} - S^*)}{\text{std}(N)}. \quad (14)$$

Note this ratio is computed with no real MR signal present. In our mock signals, EDITER yields $R \approx 30\%$ – 40% (depending on the signal), meaning 60%–70% noise reduction. The CNN achieved $R = 20\%$ – 30% , beating EDITER; however, being a calibration-based method, the CNN’s ratio is based on a null FID ($S^* = 0$).

Integration of automatic noise cancellation into the MRI setup Our current MRI console (Pure Devices) has only one analog input and thus cannot sample EMI coils simultaneously with the receive coil. We have therefore added a separate acquisition board for the EMI coils, synchronizing it with the MRI console via a digital trigger.

The MRI console samples at 125 MHz and internally demodulates to the Larmor frequency, then applies a decimation filter to produce data at the desired sampling rate (25 kHz). We perform similar operations on the acquisition board data: the board samples at 125 kHz, then we demodulate and filter the EMI data before downsampling to 25 kS/s. To avoid aliasing in the EMI data, a low-pass analog filter with cutoff $f = 62.5$ kHz is used on the EMI coil lines.

We have so far placed two EMI coils around the R_x coil. We tested several positions and orientations to optimize noise cancellation. The current optimal positions are shown in Fig. 16. We chose to proceed with EDITER since it requires no calibration data and is much faster to compute (1–3 minutes to model and correct signals, vs. ~ 10 minutes for the CNN). The ANC protocol is now fully integrated and can be used automatically with all sequences. The next section shows results of this setup on FID and 1D GRE sequences using the SQUID for amplification.

Results We first tested our ANC setup on FID sequences. Figure 17 shows the frequency-domain signal from a single shot, 10 averaged shots, and 100 averaged shots, for an example FID sequence where EMI noise came from ambient sources. No artificial EMI was added. Sequence parameters: TR = 60 ms, flip angle = 90° . The data consist of 1000 TR shots of 50 ms each, sampled at 25 kHz.

Next, we tested ANC on several 1D GRE runs using the SQUID in the detection chain. Our standard 1D GRE runs have 39 pixels over a 40 cm field of view (FOV), with console data acquired at 5 kHz over the corresponding readout duration to get 39 pixels. We found ANC performance is suboptimal with such few samples




FIG. 16. Optimal positions for the EMI coils.




FIG. 17. Original (red) vs. EDITER-corrected (purple) FID signals averaged over an increasing number of shots. The EDITER-corrected data show fewer and smaller noise spikes. The correction ratio R for this run was 0.7 (i.e., 30% noise reduction).

per shot. To remedy this, we doubled the readout duration and increased the sampling frequency to 25 kHz (same as for FIDs) to obtain a 381-pixel 1D k-space. This would increase the final image FOV if not corrected, but we avoid that by downsampling back to 5 kHz after correction, before the inverse Fourier transform.

Results of ANC on one such 1D GRE run are shown

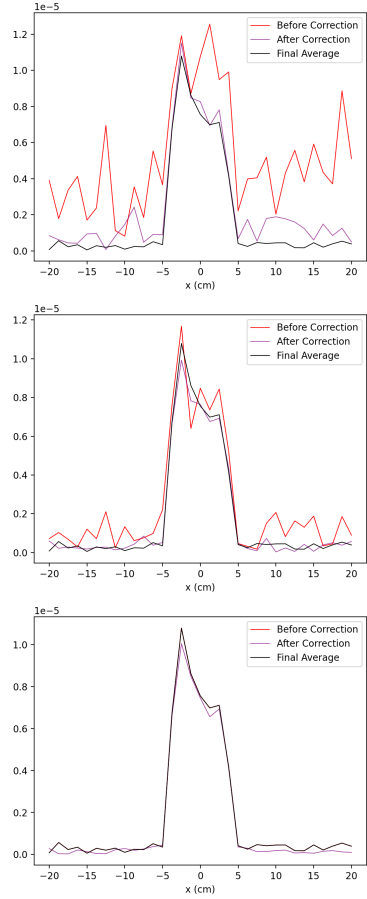


FIG. 18. Original (red) vs. EDITER-corrected (purple) 1D GRE signals averaged over an increasing number of shots. The EDITER-corrected data show fewer and smaller noise spikes.

in Fig. 18. The sequence lasted 10 minutes, using the z-gradient, with 9000 shots sampled. It can be seen that the corrected data, even after averaging very few shots, are well-denoised and match the final averaged data over all 9000 shots.

Conclusions and future work The results obtained so far with ANC on FID and 1D GRE sequences are encouraging and demonstrate the importance of methods like EDITER for reducing electromagnetic noise in ULF MRI, thereby reducing the number of averages needed for good SNR. Future work includes increasing the number and orientation of EMI coils for better sampling of all noise sources, and integrating the setup in the investigational device.

MRI reconstruction and trajectories

(Author: Guillaume Daval-Frerot.)

Sampling patterns There are three main contenders to accelerate acquisitions beyond fully sampled line-by-

line Cartesian encoding:

- Line-by-line undersampling following a Gaussian distribution, typically in the third k-space dimension (as in [34]).
- Spiral trajectories, a non-Cartesian acquisition that can be easily parameterized to match sequence specs (readout length, resolution, FOV, gradient limits, etc.). The k-space density can also be adjusted to emphasize certain regions.
- Segmented-EPI trajectories. High-field literature shows EPI can compete with far more complex non-Cartesian sampling patterns [35] like SPARKLING. At low field, readout windows will be shorter due to shorter T_1 and T_2 , but segmented EPI can still provide significant acceleration.

All these trajectories can be combined with partial Fourier acquisition, providing an extra factor of up to 2 under ideal conditions where the image phase is mostly static and low-frequency. In the absence of disturbances, k-space is Hermitian and only half of it needs to be acquired. In practice we would sample most of the central low-frequency k-space, or use navigator images to account for low-frequency phase variations.

Furthermore, 2D non-Cartesian trajectories and EPI are compatible with parallel imaging acceleration (like GRAPPA) once we have a setup with multiple receiver coils.

Our main focus at the moment is on EPI + Gaussian undersampling.

Reconstruction We currently assess performance using simulations based on the M4Raw dataset [36] acquired at 0.3 T over 183 patients with 3 contrasts (T_1 , T_2 , FLAIR). The images are modified to include various artifact sources, such as added thermal noise matching expected levels in our system, and observed B_0 inhomogeneities, etc.

Until we have a setup capable of parallel imaging and can collect deep learning datasets, reconstruction algorithms are mostly limited to partial-Fourier reconstruction (e.g., projection onto convex sets, POCS [37]) and various regularizations (e.g., compressive sensing with sparsity constraints in some wavelet domain).

The same dataset is also being used to anticipate acquisitions with multiple averages and channels, and for reconstruction using multiple contrasts [38].

Next steps Once we have target images from volunteers, we can easily explore various machine learning approaches:

- Variational networks, such as the approach Hyperfine used in a recent abstract [39]. This is the classic MR-specific deep learning approach; most state-of-the-art methods are variations of this architecture.




FIG. 19. A partial-Fourier spiral that achieves a $4\times$ acceleration over a 4 ms readout window with 10 mT/m gradients for an in-plane 2 mm target resolution, assuming no concomitant gradient issues and ideally static phase.




FIG. 20. Simulated results from the spiral acquisition of Fig. 19. For comparison, Hyperfine achieves an acceleration factor of 3.5 using deep learning ^a.

^a J. Schlemper *et al.*, Proc. ISMRM p.991 (2020).




FIG. 21. Simulation of a 2D EPI $2 \times 2 \text{ mm}^2$ acquisition of a T_2 -weighted contrast with $3\times$ acceleration, a 9 ms readout window, and only 4 mT/m gradient strength. The simulation uses the shorter T_2^* expected at low field.

- Data-driven trajectory learning, and even jointly learned trajectory + reconstruction networks. Examples include BJORK, PILOT, and PROPELLER, with a preference for BJORK [40], as it incorporates gradient hardware constraints as regularization (minimizing gradient strength and slew rate beyond just hardware limits). This helps with concomitant gradients, which become more significant when using larger or more irregular gradient patterns.
- Diffusion and score-based models [41] to provide trajectory-independent deep learning reconstruction.

Classical or quantum antennas

The core of our innovation is an ultra-sensitive magnetic detection system. Every MRI experiment needs a way to detect the signal emitted by the sample. This signal is an RF decaying pulse; for example, in a 1.5 T scanner its central frequency is about 60 MHz. At such frequencies, it is convenient to use an inductive antenna — a closed loop of wire of area S . When an electromagnetic signal passes through the loop, Faraday’s law gives:

$$Ri = -\frac{d\Phi}{dt}, \quad (15)$$

where i is the induced current in the loop, R is its resistance, and $\Phi = B_2 S$ is the flux of the signal’s magnetic field B_2 through the loop [42]. If we consider a signal at frequency $\omega_0 = \gamma B_0$, the relationship between the induced current and the polarization field B_0 is:

$$i = \frac{1}{R} \gamma B_0 B_2 S. \quad (16)$$

A sensitive antenna yields a high current i for a given field B_2 . From the above equation, inductive antennas are much more efficient at high B_0 . If we drop the field from 1.5 T to 1 mT, the antenna becomes 1500 \times less sensitive. One could try to decrease R (e.g., by cooling the coil in liquid nitrogen), but there’s no way to gain a factor of 1500 in sensitivity with such a simple change.

This is where SQUIDS come in. Unlike inductive antennas, SQUIDS have a flat frequency response over a wide bandwidth, meaning their sensitivity does not depend on frequency. This is because SQUIDS are *pure transducers* — directly sensitive to the magnetic field itself — whereas a classical antenna is sensitive to its time derivative. A direct consequence is that the lower the field, the greater the relative advantage of a SQUID over a conventional coil, making SQUIDS excellent candidates for detection at ultra-low fields.

SQUIDS offer two other advantages: extreme sensitivity and very low intrinsic noise (a few 0.1 pA/ $\sqrt{\text{Hz}}$ [43]).

These allow a detection threshold (smallest detectable signal per root bandwidth) below 1 fT/ $\sqrt{\text{Hz}}$ at very low fields. Our goal is to push even further, approaching 0.1 fT/ $\sqrt{\text{Hz}}$, which in theory (see Fig. 22) would yield detection performance equivalent to 0.2 T MRI — combined with all the advantages of ultra-low field operation.

More quantitatively, we can compare SQUIDS and classical coils as follows. For a given background noise, the detection threshold roughly scales as the ratio of detected signal to detector noise. For an inductive coil, the induced voltage $e = -d\Phi/dt$ scales as ω_0^2 (one factor of ω_0 for the field strength and another for the time derivative), whereas Johnson noise — the main noise source — scales as $\omega_0^{1/4}$ (Johnson noise $\propto \sqrt{R}$, and $R \propto \sqrt{\omega}$ due to the skin effect in conductors). Thus, an inductive coil’s detection threshold scales as $\omega_0^{-3/4}$.

For SQUIDS, the detection threshold scales simply as ω_0 . The SQUID is directly sensitive to the field, giving one factor of ω_0 , while its intrinsic noise is roughly constant over a wide frequency range. Comparing the two cases, we see a crossover in relative performance. As Fig. 22 shows, above a typical frequency (20 mT field), inductive coils have a lower detection threshold; below ~ 1 MHz (25 mT), low- T_c SQUIDS become more sensitive (their blue curve crosses the black coil curve). Low- T_c SQUIDS typically have detection thresholds of a few fT/ $\sqrt{\text{Hz}}$ (from Aimé Labbé’s thesis).

There is a wide variety of magnetic detection systems that could be promising for MRI, but SQUIDS are the most technologically mature. Even 60 years after their invention, SQUIDS remain special mainly due to their extreme sensitivity combined with a very broad bandwidth (up to tens of MHz [44]). Below we list a few other options:

- Optically pumped magnetometers (OPMs) (of the spin-exchange relaxation free (SERF) kind) based on alkali vapor cells have gained a lot of excitement in recent years, as they can reach sensitivities just above 0.1 fT/ $\sqrt{\text{Hz}}$ [45] and have been used for pre-polarized MRI at 4 mT [46]. However, they suffer notable drawbacks for NMR: most have a narrow bandwidth of at most a few hundred Hz, and they are extremely sensitive to DC inhomogeneous background fields. OPMs are also costly to produce and current industrial capacity is insufficient for large-scale deployment.
- Very low-field MRI has been demonstrated with hybrid magnetoresistive sensors by Claude Fermion’s group [47][48]. They continue to develop these sensors alongside pure ultra-low field MRI research [49].
- Nitrogen-vacancy (NV) diamond magnetometers have been explored as a technique for nanometer-scale MRI [50]. NV centers are very small room-

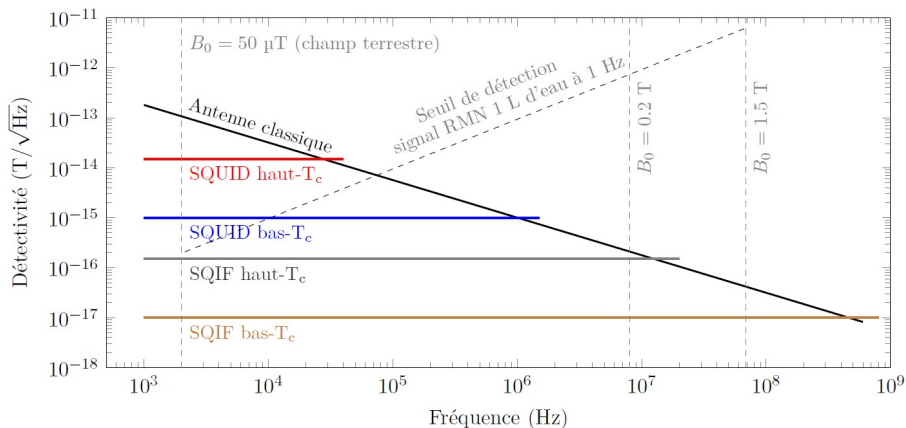


FIG. 22. Detection threshold vs. frequency for various systems. In black: inductive coils have very low detection thresholds at high frequencies (high sensitivity). For frequencies below $\simeq 1$ MHz (corresponding to ~ 25 mT), SQUIDs become more sensitive: the blue curve (low T_c single SQUID) crosses the black line of the inductive antenna. Low T_c SQUIDs have typical detection thresholds of a few $fT/\sqrt{\text{Hz}}$ (from Aimé Labbé’s thesis).

temperature magnetometers (e.g., could be placed on a catheter tip) and are ideal for detecting DC or low-frequency magnetism inside the body or near a sub-millimeter sample. However, their sensitivity threshold is at least three orders of magnitude worse than a SQUID, so they are not useful in our context of macroscopic ultra-low field MRI.

SQUIDs and quantum physics

The SQUID (superconducting quantum interference device) is a small chip that acts as an extremely sensitive magnetometer, and it lies at the heart of the ultra-sensitive detection system we are building. The first SQUID was invented in the early 1960s at Ford and Bell Labs. Below we provide some explanations of the origin of this ultra-sensitive detection phenomenon.

A typical low- T_c DC SQUID consists of a micrometer-sized loop of superconducting material (e.g., Niobium) cooled to about 4 K [51] and interrupted by two Josephson junctions — thin insulating barriers that introduce tunneling effects and allow the SQUID to act as a magnetometer.

First, let’s describe what a superconductor is and how it behaves in a magnetic field.

Current in a superconductor

The reader may recall the basics of electrical conduction in a metal like copper: driving a current I through a copper wire of resistance R produces a voltage U such that $U = RI$ (Ohm’s law). This describes how a metal

resists electron flow; resistivity arises from electrons colliding with impurities or lattice vibrations (phonons).

Certain materials (e.g., Niobium or its common alloys NbTi, Nb₃Sn) have the special property that if cooled below a critical temperature (around 10 K for Niobium), they transition to a *superconducting* state where two things happen:

- DC resistivity falls to zero.
- Ambient magnetic field is expelled from the bulk (the Meissner effect).

Such materials are called superconductors. A hand-waving explanation: in the normal state, electrons move independently and experience thermal agitation. When temperature is lowered, a very weak attractive force between electrons emerges (once thermal agitation is low enough), causing electrons to pair up. Individual electrons are fermions, but electron pairs are bosons (like photons). Bosons have the tendency to occupy the same quantum state (as in a laser where photons pile up in the same state with long-range phase coherence). Thus, all electron pairs condense into one macroscopic quantum state that is very rigid — meaning it takes a large energy to perturb it. This collective state “ignores” small defects and lattice vibrations, hence zero resistivity.

Among other consequences, Ohm’s law no longer holds: you can have a device with nonzero voltage U and nonzero current I but zero resistance. Hence, we need to find a different way to describe current in a superconductor. This was done by the London brothers in 1935 [52]. Assume that in the superconducting state, individual electrons are subject to the classical Lorentz force, $\vec{F} = -e(\vec{E} + \vec{v} \times \vec{B})$, where \vec{E} is the electric field, \vec{B} is the magnetic field, $-e$ is the electron charge, and \vec{v} the

electron's velocity. For nonrelativistic electron velocities ($v \ll c$, where c is the speed of light in vacuum) we can neglect the magnetic term. Applying Newton's second law to this electron yields:

$$m \frac{d\vec{v}}{dt} = -e\vec{E}. \quad (17)$$

Now introduce the volumetric supercurrent density [53] $\vec{j} = n_S e \vec{v}$, where n_S is a phenomenological parameter defined by the London brothers as the density of electrons participating in the superconducting state. This gives the first London equation:

$$\frac{d\vec{j}}{dt} = \frac{n_S e^2}{m} \vec{E}. \quad (18)$$

This tells us about the time variation of the current. Now, let's examine the structure of this current. Taking the curl of the first London equation and using Maxwell's equation $\nabla \times \vec{E} = -\partial \vec{B} / \partial t$, we get:

$$\frac{\partial}{\partial t} \left(\nabla \times \vec{j} + \frac{n_S e^2}{m} \vec{B} \right) = 0. \quad (19)$$

A superconductor has two fundamental properties: zero electrical resistance, and another property called the *Meissner effect*, which is that the magnetic field cannot penetrate inside the bulk of the material ($\vec{B} = 0$). Because of the Meissner effect, we immediately see that a constant nonzero solution to the above equation would be unphysical. Hence:

$$\frac{\partial}{\partial t} (\nabla \times \vec{j}) = 0 \implies \nabla \times \vec{j} = -\frac{n_S e^2}{m} \vec{B}. \quad (20)$$

This is the second London equation.

At this point, it is useful to introduce the vector potential \vec{A} through $\vec{B} = \nabla \times \vec{A}$. As with any potential function (like potential energy in classical mechanics), it is defined only up to a certain reference. Fixing this reference is called choosing a *gauge*. In the case of the vector potential \vec{A} , due to Maxwell's equations, it is defined up to a vector field with zero curl. We choose the Coulomb gauge, in which $\nabla \cdot \vec{A} = 0$. In this gauge [54], the two London equations are nicely summed up into one:

$$\vec{j} = -\frac{n_S e^2}{m} \vec{A}. \quad (21)$$

U(1) symmetry, flux quantization

As discussed previously, electrons in a superconductor all bind into one big macroscopic superconducting state described by a wavefunction $\Psi(r)$ that depends on position r in the material (in our case, the SQUID loop). This wavefunction has the property that it is defined only up to a constant phase: the transformation $\Psi \rightarrow \Psi e^{i\chi}$,

where χ is a real constant, leaves the quantum state invariant. The group associated with this transformation is called $U(1)$ and is a global symmetry [55] of our system.

In most microscopic quantum theories, when you have a global symmetry of your system, it is tempting to assume it is also *local*, meaning the phase in the symmetry group depends on position. This "trick" is far more profound than it sounds, as it unveils features of the system that were otherwise hidden in the global symmetry. Such local theories are called *gauge theories*, and are the backbone of most modern microscopic theories: quantum electrodynamics (describing light-matter interaction), or quantum chromodynamics/electroweak theory (components of the Standard Model of particle physics). A famous manifestation of gauge theory is the emergence of particle masses through the coupling to a massive scalar field, the Higgs boson. But let's get back to the matter at hand.

Assume that the transformation $\Psi(r) \rightarrow \Psi e^{i\chi(r)}$ becomes local. The Hamiltonian of the system has a kinetic energy term given by $\frac{(\vec{p} - e\vec{A})^2}{2m}$, where the momentum operator $\vec{p} = i\hbar\nabla$. To maintain invariance under the local symmetry $\Psi(r) \rightarrow \Psi e^{i\chi(r)}$, the reader is invited to check that the vector potential field must transform as:

$$\vec{A}(r) \rightarrow \vec{A}'(r) = \vec{A}(r) + \frac{\hbar}{e} \nabla \chi(r). \quad (22)$$

Now define $\phi(r) = 2\chi(r)$ [56].

To respect gauge invariance, the London equation in (21) must be rewritten as:

$$\vec{j}(r) = \frac{en_S}{2m} \left(\hbar \nabla \phi(r) - 2e\vec{A}(r) \right). \quad (23)$$

Now we can get to the core principle governing a SQUID: the phenomenon of flux quantization.

A SQUID [57] is a small loop made of superconducting material interrupted by two Josephson junctions, which are small barriers splitting the loop into two parts.

Assume we have a magnetic field B perpendicular to the loop. The flux of the magnetic field through the SQUID is $\Phi = \iint_S \vec{B} \cdot d\vec{S}$, where S is the surface of the disk delimited by the SQUID loop. A supercurrent $\vec{j}(r)$ is induced in the loop by the magnetic field. Let us write the circulation of this supercurrent around the loop:

$$\oint \vec{j}(r) \cdot d\vec{l} = \frac{en_S}{2m} \left(\hbar \oint \nabla \phi(r) \cdot d\vec{l} - 2e \oint \vec{A}(r) \cdot d\vec{l} \right). \quad (24)$$

Let's analyze these two terms one by one. The first one gives:

$$\oint \nabla \phi(r) \cdot d\vec{l} = \phi(r)_{\text{after one turn}} - \phi(r). \quad (25)$$

Phase is defined up to an integer number of times 2π :

$$\oint \nabla \phi(r) \cdot d\vec{l} = 2n\pi. \quad (26)$$

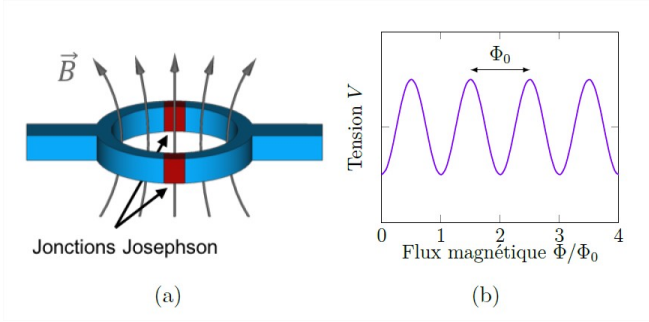


FIG. 23. (a) Schematic drawing of a SQUID. (b) Voltage–flux characteristic of a SQUID. When a magnetic flux Φ threads the SQUID, a voltage drop V appears. It varies periodically with flux, period equal to one flux quantum Φ_0 . Measuring this voltage drop $V(\Phi)$ allows one to infer the flux Φ , and hence the magnetic field B we want to measure.

For the second term, we use Stokes’s theorem:

$$\oint \vec{A}(r) \cdot d\vec{l} = \iint_S (\nabla \times \vec{A}(r)) \cdot d\vec{S} = \Phi, \quad (27)$$

yielding:

$$\oint \vec{j}(r) \cdot d\vec{l} = \frac{en_s}{2m} (2\hbar\pi n - 2e\Phi). \quad (28)$$

In the superconducting state, supercurrents flow only at the surface of the material. This can be seen as a consequence of the Meissner effect: magnetic field cannot penetrate the bulk of the material. With this in mind, let’s choose a path C for the integration that passes through the bulk of the SQUID loop, where $\vec{j}_s = 0$ everywhere. We obtain:

$$\Phi = n \frac{\hbar}{2e}, \quad (29)$$

where n is an integer. In other words, the magnetic flux threading the SQUID’s loop is quantized: it can only take values that are multiples of the flux quantum $\Phi_0 = \frac{\hbar}{2e}$.

What we learn from this is that when a magnetic field threads the SQUID loop with flux Φ , the SQUID will respond by generating a small current I . This current creates a small counter-flux $\delta\Phi$ that will re-establish the total flux to the nearest integer multiple of Φ_0 :

$$\Phi_{\text{tot}} = \Phi - \delta\Phi = n\Phi_0. \quad (30)$$

To exploit this phenomenon, we will add another ingredient to our superconducting loop: two Josephson junctions in parallel, as described in Fig. 23. We won’t go too much into the details of the physics of Josephson junctions as it would be a bit too heavy for this introductory white paper. A very good description of Josephson equations in a SQUID is given in the PhD thesis of Nazim Lechéa (Part III.1) [58], if you read French. Also,

another excellent reference is the SQUID Handbook by John Clarke (Part 2.1) [59]. Below we give the main ideas behind how a SQUID reads a magnetic field.

In the presence of the two junctions, the current circulating in the loop is given by:

$$I(\delta) = I_0 \sin(\delta), \quad (31)$$

where I_0 is the junction critical current, and δ represents the total dephasing of the system’s wavefunction across the junctions. It is given by $\delta = \Delta\phi_b + \Delta\Phi_{\text{mag}}$, where $\Delta\phi_b$ is the ”bare” dephasing of the wavefunction without the presence of the magnetic field, and $\Delta\Phi_{\text{mag}} = 2\pi\Phi_a/\Phi_0$ the dephasing due to the magnetic flux Φ_a threading the loop. This equation is the first Josephson equation and describes the system well for small current intensities $I < I_0$. What’s more interesting is to pre-bias the system to a current close to I_0 . When the current generated by the magnetic flux goes above this limit, a tension U_j appears across one junction, given by:

$$U_j = \frac{\Phi_0}{2\pi} \frac{d\delta(t)}{dt}, \quad (32)$$

which is the second Josephson equation. Solving this pair of equations for the whole system, we get that the tension $U(\Phi_a)$ across the SQUID varies periodically with the applied magnetic flux Φ_a in the form:

$$U(\Phi_a) = U_0 \sin\left(\frac{2\pi\Phi_a}{\Phi_0}\right). \quad (33)$$

For small flux amplitudes (which the SQUID is designed to measure) near a point of maximum slope, the variation of tension ΔU is linearly related to the variation of magnetic flux $\Delta\Phi$ via:

$$\Delta U = 2\pi U_0 \frac{\Delta\Phi}{\Phi_0}. \quad (34)$$

For very low fields in the pT range, this tension is on the order of a few μV . It then gets amplified by low-temperature preamps to reach a few V in room-temperature electronics. This makes SQUIDS extremely sensitive magnetometers, with noise levels in the $\text{fT}/\sqrt{\text{Hz}}$ range.

Elements of a SQUID-based antenna

A typical SQUID-based MRI detection system always consists of a few core parts:

- A flux concentrator, whose job is to concentrate the magnetic signal towards the SQUID.
- One or several SQUIDS, which read the incident signal.
- A low-temperature preamplification system.




FIG. 24. A simplified view of a typical SQUID detection system for an MRI experiment. On the left is the primary antenna, a resonant saddle coil in this example, that captures the signal and sends it to the SQUID loop (in the middle of the picture, the circle with two crosses representing the Josephson junctions). The SQUID reads the signal, and its output is amplified by a low-temperature cryogenic preamplification system. Then, the signal is sent to room-temperature electronics which ensure amplification, feedback, and analog-to-digital conversion.

- Room-temperature electronics to provide amplification, feedback, and analog-to-digital conversion.

I will say a few words about the first element, the flux concentrator, as it is highly critical regarding the sensitivity of the whole system. A SQUID can be imagined as a small ring made of superconducting material, interrupted by two "barriers" (Josephson junctions), with diameter on the order of μm . Because of its small size, the SQUID intercepts only a tiny fraction of the sample's signal. A common way to circumvent this issue is to couple the SQUID to a *flux concentrator*. It is just a common loop antenna, called the *pickup coil*, in series with an *input coil*, which is a small coil that sends the captured flux through the SQUID. You can think of it as the pickup coil capturing the signal and sending it to the SQUID that reads it.

A very thorough explanation of the different types of flux concentrators one can use has been done by Fagaly [60]. There are two important parameters to choose when designing the pickup coil: its geometry, and the nature of the wire you're going to use. Usually, coils fall into two geometrical categories: surface or volume coils. Surface coils are often smaller and allow minimal intrinsic noise, and also allow designing *gradiometric* geometries that reject far-field noise. The trick is to split the antenna into two parts that run against each other. Objects close to the antenna send out signals that will get detected, but noise from distant sources induces equal signals into the two parts of the antenna. As a consequence, signals in the two parts cancel each other, and the resulting signal averages to zero.

Volume coils have been less explored in the context of SQUID MRI but they are standard in the case of classical inductive MRI, where saddle-coil or birdcage geometries are often used [61].

Regarding wire choice, you have the choice between a classical ohmic conductor such as copper or a low-temperature superconductor, Niobium (or Niobium-Titanium alloy) wire being a standard. Copper wire has the advantage of being able to operate at any temperature, at the expense of producing more Johnson-Nyquist noise. This voltage noise goes as $\Delta u = \sqrt{4k_B T R \Delta f}$,

where k_B is Boltzmann's constant, T is temperature, R is the wire's DC resistance and Δf is the bandwidth over which you're measuring the noise. To minimize this noise, it is a good idea to choose high-quality OFHC copper that has a very low resistivity, and to cool down the antenna (for instance in liquid nitrogen). Niobium wire, on the other hand, is a superconductor and has to be cooled below 10 K to be used. Because of this, a niobium pickup coil has to stay at the heart of the cryostat, far away from the object you need to measure, and only a handful of surface geometries can be used. Of course, the Johnson noise is going to be a lot less in a Niobium wire. Hence, choosing Copper or Niobium wire will be a trade-off between flexibility on the geometry of the coil, and minimizing the intrinsic Johnson noise of the antenna.

Until now, most low- T_c SQUID-detected MRI systems have used superconducting second-order axial gradiometers, see e.g. Zotev 2007 [62] or Penanen 2014 [63]. Some high- T_c experiments have used ohmic volume antennas, such as the device of Chen et al. 2011 [64]. At Chipiron, we believe that the geometry of this primary antenna is the key to increase the SNR enough to be able to perform SQUID MRI in a clinical setting.

Building and using a SQUID detection system

The reason we are using SQUIDs is, of course, because they have exquisite sensitivity. This also means that installing and running a SQUID detection system requires many precautions to properly thermalize the system and avoid picking up interference.

All the leads to the SQUID, including the pickup coil and the output cables, have to be properly heat sunk. Unlike most SQUID experiments that are performed inside a wet cryostat using liquid He, we use a pulse-tube cryocooler and the SQUID is cooled on a cold metallic plate in vacuum. All cables and the SQUID have to be cooled with a thermal metallic link to the plate. This has to be done extremely carefully to avoid shorts, ground loops, and picking up interference. This includes careful soldering, proper shielding of all cables, and the use of twisted pairs to reduce noise to the minimum. Because




FIG. 19. (a) Magnetometer; (b) first derivative axial gradiometer; (c) first derivative planar gradiometer; (d) second derivative axial gradiometer; (e) second derivative asymmetric axial gradiometer; (f) first derivative radial gradiometer.

FIG. 25. A few examples of geometries for surface antennas. Apart from antenna (a), all other geometries are of gradiometric nature: the several coil turns compensate each other. If you call b the distance between two of the loops, signals from sources at distance b or less will get detected. On the contrary, if the emitting object lies at a distance $\gg b$, the different coils compensate each other and the resulting signal is zero. Sketches (b)–(f) present several alternatives of gradiometers of different orders and different geometries. The higher the order, the more effective you are at rejecting far-field noise but the less sensitive you become. (From Fagaly et al., 2006.)




FIG. 26. A SQUID magnetometer, model SQ2600 of STAR Cryoelectronics mounted with readout electronics and a Niobium shield assembly (shield open in the picture). The SQUID chip is protected under the resin layer on the left side of the picture. The input coil can be accessed through the two Niobium pads; leads are introduced in the shield through the thin metallic tube on the top. We use a copper braid as a heat sink to properly thermalize the chip by connecting it to the cold plate of the cryostat.




FIG. 27. Several washer design SQUID chips under the microscope. The two big squares on the chip are the input coil; the SQUID's Josephson junctions can be seen between the two coils in the center. The different elements (feedback coil, SQUID output, input coil) are reached through the gold pads on the edges of the chip.

everything is metallic, it is very tricky to keep the pickup coil and the SQUID electrically floating (to avoid ground loops) while ensuring good thermalization. On top of this, it is useful to use cryogenic transformers for ground isolation of the SQUID. Analog filtering on the input line is also important for two reasons: to improve the signal-to-noise ratio, and to prevent high-frequency signals from destabilizing the SQUID's feedback electronics. Typical noise picked up by the cables comes from instruments in the lab, elevators, and passing cars or trains. In our case, the biggest source of noise was *microphony* from the pulse-tube cryocooler, which comes from induced current inside vibrating cables in an inhomogeneous background magnetic field. The cryocooler works with a cycle of compression-decompression of helium gas. In comparison with other similar cryocoolers like Gifford-McMahon, pulse tubes produce a lot less mechanical vibrations, because there are no moving metallic parts in the cold head. Still, there remain some unavoidable residual vibrations that propagate through the rigid parts of the cryostat and the cables that are bound to it. There are three ways to curb microphonic noise:

- Using shielded cables, to limit the interaction with background magnetic fields.
- Using cables wound in twisted pairs. Locally, the currents created in each cable cancel each other, which limits the noise.
- Reducing mechanical vibrations to the minimum. For this last point, we modified our cryostat to tie the cold plates to the pulse-tube with copper braids. This gives some flexibility to the structure and kills most of the very low frequency noise below 1 kHz.

Once the system is all set up with protected and well heat-sunk leads, we can start some magnetic detection measurements. First you start by directly injecting some

current in the SQUID input, either through the input coil via a low-noise stabilized current source or via the feedback coil. This does not require a pickup coil to begin with. This first step is used to set the parameters of the readout electronics, mainly the gain of the pre-amplifier and flux lock loop. This controls the detection bandwidth (which we want to keep up to 100 kHz in our case) and the dynamic range (less important because we measure very tiny signals on the order of 1 pT). One must be very cautious with the amplitude of the current — anything above 100 μA can destroy the chip by creating a short between one of the input coils and the SQUID.

Now, the next step is to perform a magnetic detection experiment with a pickup coil, through radiative coupling. To start, it is easier to use a reduced-size pickup coil wound from superconducting Niobium wire, as in most SQUID MRI experiments. With this configuration, we can test the effect of the geometry in a very controlled environment, where the very low temperatures suppress most of the Johnson noise. Also, we can enclose the detection system in a Pb box, lead being superconducting below 7 K. Because of the Meissner effect, superconducting shields are extremely effective for magnetic shielding, much more than typical room-temperature magnetic materials. Then, once environmental and intrinsic noise has been reduced to a minimum, one can place a source of signal inside the shield, namely a small copper loop carrying a small alternating current at a given frequency. As SQUIDs are so sensitive, it is extremely important to reduce the noise to a minimum and to use very small signals for testing, as anything too intense will quickly put the flux-lock loop out of its operating range, and more generally saturate the detection system. With this configuration, we can safely compare the responses of different sizes and geometries of flux concentrators with the same source and background noise, yielding comparisons of signal-to-noise ratios.

To test our technology in a more realistic environment, we need to perform the same detection experiment with a full-scale copper antenna at room temperature. In our eventual product, the antenna will have to enclose the body part of interest; in our case, a part of the human body. With this configuration the antenna cannot be confined to the depths of the 4 K cryostat. Our simulations show that with a cylindrical antenna of typical length 30 cm and diameter 20 cm, it is realistic to cool the antenna to $\sim 40\text{--}50$ K while keeping the cryogenic requirements low enough to design a light, open cryostat adaptable to many parts of the body. This means that Niobium superconducting antennas are off the table, and that we will rather resort to high-quality copper antennas, or, if our trials are successful, medium- T_c superconducting magnesium diboride (MgB_2) antennas (currently being tested). We built this light cryostat that is currently under test.

As pointed out at the beginning of this part, running

leads from the 4 K SQUID to the room-temperature coil raises a number of challenges: proper cable thermalization, ground isolation, mechanical vibrations that turn into noise via microphony. We are currently solving these problems with cryotransformers, careful wiring, and mechanical vibration damping. To deal with the extreme sensitivity of the SQUID, in the fully cold 4 K case we used superconducting lead shields to reduce environmental noise and allow us to use extremely weak signals. At room temperature, obviously, it is hopeless to use superconducting shields, hence active shielding strategies like the one described in Yilong Liu et al. [65] will be used.

* Corresponding author: dimitri@chipiron.co

- [1] Srinivas *et al.*, Magn. Reson. Med. **87**(2), 614–628 (2022)..
- [2] Yilong Liu *et al.*, Nature Commun. **12**, 7238 (2021)..
- [3] P. W. Adamson *et al.*, Magn. Reson. Med. 1–14 (2025).
- [4] M. H. Mazurek *et al.*, Nat. Commun. **12**, 5119 (2021).
- [5] <https://investors.hyperfine.io/static-files/dbda7edd-40f1-4a26-b1f7-3eb5c1f9c988>.
- [6] A. J. Sorby-Adams *et al.*, Nat. Commun. **15**, 10488 (2024).
- [7] S. Busch, PhD dissertation, UC Berkeley, 2011.
- [8] Here $f_0 = \omega_0/2\pi$ is the frequency in Hz corresponding to angular frequency ω_0 .
- [9] We neglect self-inductance for simplicity.
- [10] We prefer the term *coherent sensors* since these sensors exploit long-range coherence, but the term.
- [11] This configuration is known as a transimpedance amplifier.
- [12] Andrew Webb, a renowned low-field MRI researcher, famously said that low-field MRI is not a replacement for high-field, but a replacement for no field.
- [13] M. Bödenler *et al.*, Magn. Reson. Med. **86**, 2049–2063 (2021).
- [14] S. Busch *et al.*, Magn. Reson. Med. **67**(4), 1138–1145 (2012).
- [15] R. Ayde *et al.*, NMR Biomed. **38**(1): e5268 (2025).
- [16] A. J. Sorby-Adams *et al.*, Nat. Commun. **15**, 10488 (2024).
- [17] Y. Zhao *et al.*, Science **384**, 6696 (2024).
- [18] L. Weizman *et al.*, Med. Phys. **42**(9), 5195–5208 (2015).
- [19] M. L. de Leeuw den Bouter *et al.*, Sci. Rep. **12**, 6362 (2022).
- [20] A. Salehi *et al.*, J. Mag. Res. **370**, 107812 (2025).
- [21] M. U. Akbar *et al.*, Scientific Data **11**, 259 (2024).
- [22] J. E. Iglesias *et al.*, NeuroImage **237**, 118206 (2021).
- [23] M. Lyu *et al.*, Sci. Data **10**, 264 (2023).
- [24] P. A. Bottomley *et al.*, Med. Phys. **11**(4) (1984).
- [25] M. L. de Leeuw den Bouter *et al.*, Sci. Rep. **12**, 6362 (2022).
- [26] M. Zach *et al.*, IEEE Trans. Med. Imaging **42**(12) (2023).
- [27] Srinivas *et al.*, Magn. Reson. Med. **87**(2), 614–628 (2022).
- [28] Yilong Liu *et al.*, Nat. Commun. **12**, 7238 (2021).
- [29] Su *et al.*, IEEE Trans. Med. Imaging **41**(5), 1007–1016 (2022).
- [30] Parsa *et al.*, J. Magn. Reson. **346**, 107355 (2023).
- [31] Srinivas *et al.*, Magn. Reson. Med. **87**(2), 614–628 (2022).

- [32] Srinivas *et al.*, Magn. Reson. Med. **87**(2), 614–628 (2022).
- [33] Parsa *et al.*, J. Magn. Reson. **346**, 107355 (2023).
- [34] This technique maps measured NMR signals to simulations. See Dan Ma *et al.*, Nature **495**, 187–192 (2013).
- [35] P. Sati *et al.*, Mult. Scler. J. **20**(11), 1464–1470 (2014).
- [36] M. Lyu *et al.*, Sci. Data **10**, 264 (2023).
- [37] G. McGibney *et al.*, Magn. Reson. Med. **30**(1), 51–59 (1993).
- [38] J. Huang *et al.*, Magn. Reson. Imaging **32**(10), 1344–1352 (2014).
- [39] Z. Ramzi *et al.*, Appl. Sci. **10**(5), 1816 (2020).
- [40] G. Wang *et al.*, IEEE Trans. Med. Imaging **41**(9), 2318–2330 (2022).
- [41] M. Zach *et al.*, arXiv:2210.13834v3 (2023).
- [42] We neglected self-inductance for simplicity.
- [43] From the STAR Cryoelectronics SQUID datasheet.
- [44] D. Drung *et al.*, DC SQUID readout electronics with up to 100 MHz closed-loop bandwidth, IEEE Trans. Appl. Supercond. **15**(2), 777–780 (2005).
- [45] Dang *et al.*, Appl. Phys. Lett. **97**, 151110 (2010).
- [46] Savukov *et al.*, Appl. Phys. Lett. **103**, 043703 (2013).
- [47] Herreros *et al.*, Rev. Sci. Instrum. **84**, 095116 (2013).
- [48] A. Doll *et al.*, Proc. SPIE **11090**, Spintronics XII, 110903M (2019).
- [49] S. Lecurieux-Lafayette, PhD Thesis, 2021, Université Paris-Saclay, <https://www.theses.fr/2021UPASP147>.
- [50] P. Hemmer *et al.*, Science **339**(6119), 529–530 (2013).
- [51] Around -270°C .
- [52] To avoid excessive technicality, we stick to a semiclassical discussion. Of course, as mentioned, electrons do not move freely in a superconductor, but this simple model captures the right phenomenology for current–voltage behavior.
- [53] Recall that electrical current I is the flux of current density \vec{j} through cross-sectional area S of the conductor.
- [54] It is interesting to note that the current density is a well-defined (gauge-invariant) quantity while the vector potential is not (only defined up to a curl-free vector). The above equation only remains true in the Coulomb gauge, and remains approximately true for slowly spatially varying fields.
- [55] A symmetry in physics is a group of transformations that leaves the system invariant.
- [56] Without going too deep into details, the phase $\phi(r)$ represents the.
- [57] We are interested in DC SQUIDS; RF SQUIDS are not discussed here.
- [58] <https://www.theses.fr/175465479>.
- [59] J. Clarke and A. I. Braginski, *The SQUID Handbook*, vol. 1, Wiley-VCH (2004).
- [60] R. L. Fagaly, Rev. Sci. Instrum. **77**, 101101 (2006).
- [61] D. Gruber *et al.*, J. Magn. Reson. Imaging **48**(3), 590–604 (2018).
- [62] V. S. Zotev *et al.*, Supercond. Sci. Technol. **20**, S367 (2007).
- [63] M. Penanen *et al.*, Rev. Sci. Instrum. **85**, 094302 (2014).
- [64] H.-H. Chen *et al.*, J. Appl. Phys. **110**, 093903 (2011).
- [65] Yilong Liu *et al.*, Nat. Commun. **12**, 7238 (2021).