Electrical Impedance Tomography: Computation and Applications

Part 4: From the Continuum Model to EIT Data

Melody Alsaker

Jyväskylä Summer School 2026

August 2026

In the previous section, we arrived at the mental model

$$\text{conductivity } \sigma(x) + \text{applied boundary current} \longrightarrow \text{boundary voltages.}$$

But what, precisely, do we mean by an *applied boundary current* or a *boundary voltage*?

We will begin with an idealized continuum model, then move toward the electrode measurements made by real EIT systems.

The Continuum Model

The continuum model assumes ideal access to the entire boundary.

- The target is represented by a domain $\Omega \subset \mathbb{R}^2$ with boundary $\partial\Omega$.
- The applied current density is represented by a function

$$j : \partial\Omega \to \mathbb{R}.$$

- The resulting boundary voltage is represented by a function

$$v : \partial\Omega \to \mathbb{R}.$$

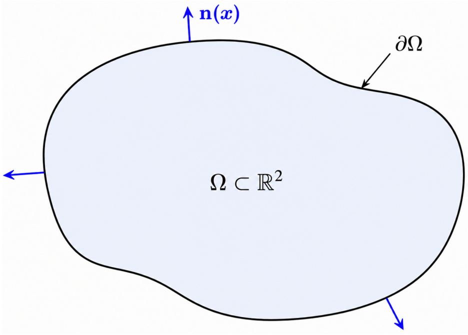

Thus, the boundary data are functions of position, not finite lists of electrode measurements.

Current at the Boundary

Let $\mathbf{J}(x)$ denote the current density inside Ω, and let $\mathbf{n}(x)$ be the outward unit normal vector at the boundary.

Only the component of $\mathbf{J}$ normal to the boundary describes current crossing the boundary:

$$\mathbf{J} \cdot \mathbf{n} = \text{signed outward current flux density}$$

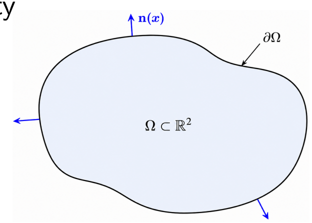

- $\mathbf{J} \cdot \mathbf{n} > 0$: current is leaving Ω,
- $\mathbf{J} \cdot \mathbf{n} < 0$: current is entering Ω,
- $\mathbf{J} \cdot \mathbf{n} = 0$: current is locally tangent to the boundary.

We define the applied boundary current density to be positive inward:

$$\boxed{j = -\mathbf{J} \cdot \mathbf{n}}$$

The Boundary Current Condition

Recall Ohm's law: $\mathbf{J} = -\sigma \nabla u$.

Since the normal derivative is given by $\frac{\partial u}{\partial \mathbf{n}} = \nabla u \cdot \mathbf{n}$,
we have

$$\mathbf{J} \cdot \mathbf{n} = (-\sigma \nabla u) \cdot \mathbf{n} = -\sigma \frac{\partial u}{\partial \mathbf{n}}.$$

Therefore,

Applied Boundary Current Density

$$\boxed{j = -\mathbf{J} \cdot \mathbf{n} = \sigma \frac{\partial u}{\partial \mathbf{n}} \qquad \text{on } \partial\Omega.}$$

Applying a current pattern means prescribing a Neumann boundary condition for u.

Current Patterns Must Balance

There are no current sources or sinks inside the target, so

$$\nabla \cdot \mathbf{J} = 0 \qquad \text{in } \Omega.$$

The divergence theorem then gives

$$0 = \int_{\Omega} \nabla \cdot \mathbf{J}\, dx = \int_{\partial\Omega} \mathbf{J} \cdot \mathbf{n}\, ds = -\int_{\partial\Omega} j\, ds.$$

Therefore, any physically admissible current pattern must satisfy

$$\boxed{\int_{\partial\Omega} j\, ds = 0.}$$

The total current entering the target must equal the total current leaving it.

Boundary Voltage

The electrical potential throughout the target is denoted by $u(x)$.

Its value on the boundary is the boundary voltage data:

Boundary Voltage

$$\boxed{v = u|_{\partial\Omega}}$$

The voltage data represents a Dirichlet boundary condition for u.

Strictly speaking, only voltage *differences* are physically meaningful, since u and $u + C$ produce exactly the same electric field and current density.

We can choose a reference by requiring, for example,

$$\int_{\partial\Omega} v \, ds = 0.$$

The Continuum Forward Problem

Suppose we know the conductivity $\sigma(x) > 0$ and choose an applied current pattern j satisfying

$$\int_{\partial\Omega} j \, ds = 0.$$

We solve

$$\begin{cases} \nabla \cdot (\sigma \nabla u) = 0 & \text{in } \Omega, \\ \sigma \dfrac{\partial u}{\partial \mathbf{n}} = j & \text{on } \partial\Omega, \\ \displaystyle\int_{\partial\Omega} u \, ds = 0 & \text{(reference choice).} \end{cases}$$

The predicted boundary voltage is then

$$\boxed{v = u\big|_{\partial\Omega}.}$$

From One Experiment to a Boundary Map

For one applied current pattern j, we obtain one boundary voltage response v.

Repeating this with many different current patterns probes the conductivity from many different directions.

The continuum *Neumann-to-Dirichlet map* packages these responses:

$$\boxed{\mathcal{R}_\sigma : j \longmapsto v.}$$

$$\text{applied boundary current} \quad \longrightarrow \quad \text{measured boundary voltage}$$

The full continuum model assumes that this response is known for every admissible current pattern.

Neumann-to-Dirichlet or Dirichlet-to-Neumann?

In most real-world EIT setups, current is applied and the resulting voltage is measured.

In the ideal continuum model, this gives the *Neumann-to-Dirichlet* (ND) ordering:

$$\mathcal{R}_\sigma : \underbrace{j}_{\text{current}} \longmapsto \underbrace{v}_{\text{voltage}} .$$

But classical inverse problem theory often uses the *Dirichlet-to-Neumann* (DN) map:

$$\Lambda_\sigma : \underbrace{v}_{\text{voltage}} \longmapsto \underbrace{j}_{\text{current}} .$$

After restricting to balanced currents and fixing a voltage reference, it turns out that

$$\boxed{\mathcal{R}_\sigma = \Lambda_\sigma^{-1}.}$$

Thus, the two maps contain the same idealized boundary information; the choice depends on whether we emphasize measurement practice or inverse-problem theory.

The Continuum Inverse Problem

Inverse Problem

Given knowledge of the ND map $\mathcal{R}_\sigma$ (or, equivalently, the DN map Λ_σ), recover the conductivity distribution $\sigma(x), \quad x \in \Omega$.

But the continuum model has made some strong idealizations:

- currents and voltages are functions on the entire boundary,
- the entire boundary is accessible,
- electrodes and contact effects are ignored.

Next, we replace this idealized boundary access with a finite collection of electrodes.

From the Continuum Model to Electrodes

The continuum model assumed that currents and voltages were known as functions on the entire boundary:

$$j(x) \longmapsto v(x), \qquad x \in \partial\Omega.$$

A real EIT system has only a finite collection of electrodes $e_1, e_2, \ldots, e_L \subset \partial\Omega$.

Input

Total electrode currents

$$\mathbf{I} = [I_1, \ldots, I_L]^\top$$

Output

Electrode voltages

$$\mathbf{V} = [V_1, \ldots, V_L]^\top$$

We will first examine what a real EIT instrument applies and measures. Later, we will develop a mathematical model that predicts these measurements.

Our Electrode Geometry

As we saw in the applications section, electrodes can be arranged in many different ways on the surface of a target.

But we will narrow our focus to a common and mathematically convenient geometry:

- The electrodes lie in a single plane.
- They are arranged in a ring around the target.
- We model the resulting two-dimensional cross-section.

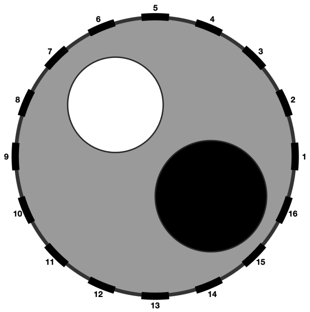

For now, let's also imagine a circular domain with $L = 16$ electrodes, as shown.

Our example domain has one conductive inclusion (white) and one resistive inclusion (black).

Example: Adjacent Current Injection

In principle, current may be injected and withdrawn on any of the electrodes, as long as the net inflow is 0.

As a starting point, let's consider the commonly used *adjacent drive* strategy:

- Electrode e_ℓ injects current $+I$.
- Electrode $e_{\ell+1}$ withdraws current $-I$.
- All remaining electrodes have zero applied net current.

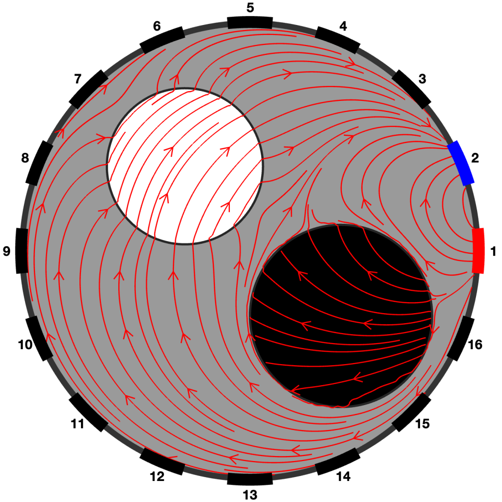

The simulated streamlines of the resulting current density field $\mathbf{J}$ are plotted.

The Induced Potential Field

Since $\mathbf{J} = -\sigma \nabla u$, the current injection also induces an internal potential field u.

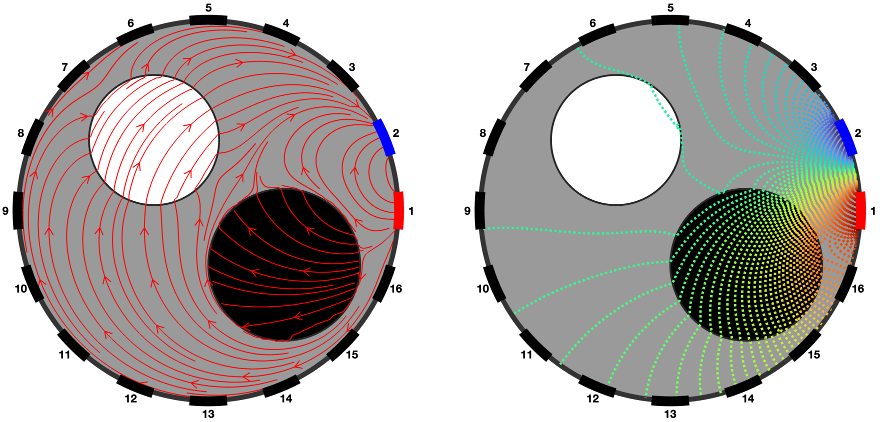

The internal field u gives rise to measurable potential differences at the electrodes.

How Do Voltage Measurements Work?

We assume each electrode e_ℓ is a very good conductor. Therefore, each electrode itself has a single voltage:

$$\boxed{V_\ell = \text{voltage of electrode } e_\ell.}$$

An instrument may record:

- electrode voltages relative to a chosen reference, or
- voltage differences $V_p - V_q$ between pairs of electrodes.

Because only voltage differences matter, we can choose the normalization

$$\sum_{\ell=1}^{L} V_\ell = 0.$$

After choosing a reference, we represent the voltage data by one value V_ℓ per electrode.

Resulting Voltage Measurement

So for our case of adjacent drive electrodes with $\ell = 1$, we obtain a current vector

$$\mathbf{I}^{(1)} = [I, -I, 0, 0, \ldots, 0]^\top.$$

Assume we obtain one voltage measurement for every electrode, including the drive electrodes.

This yields a corresponding voltage vector

$$\mathbf{V}^{(1)} = [V_1^{(1)}, V_2^{(1)}, \ldots, V_{16}^{(1)}]^\top.$$

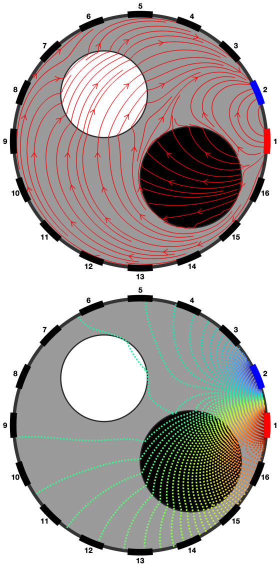

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

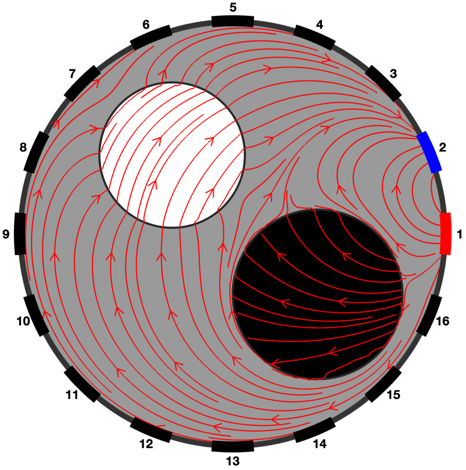

Inject $\mathbf{I}^{(1)}$, Measure $\mathbf{V}^{(1)}$.

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

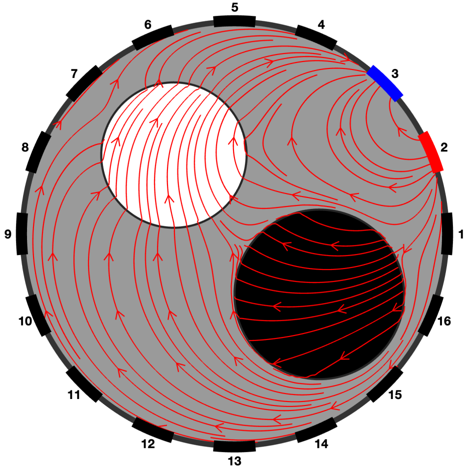

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

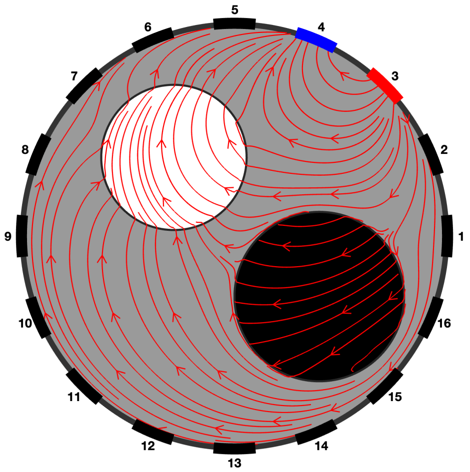

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

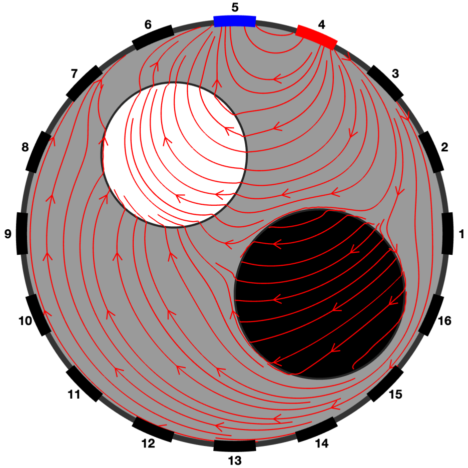

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

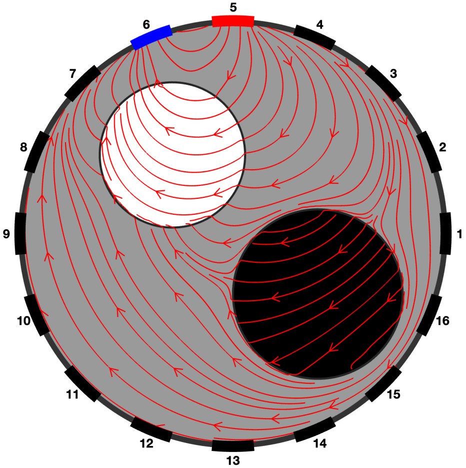

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

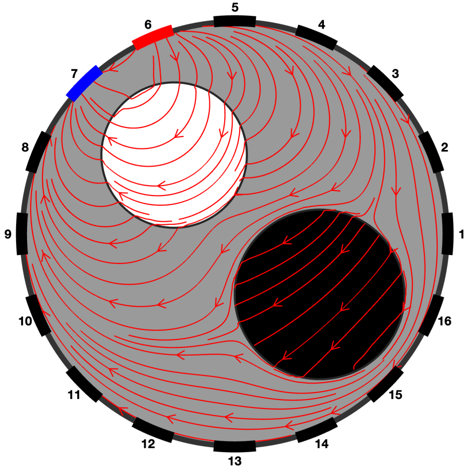

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

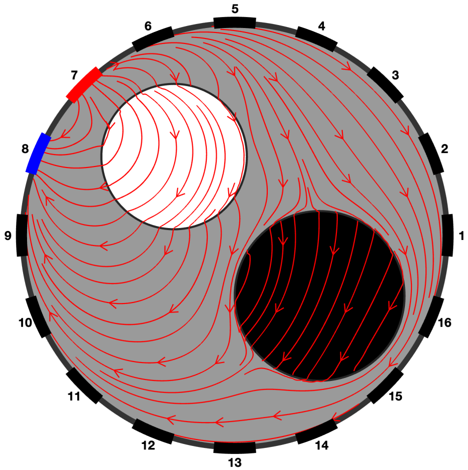

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

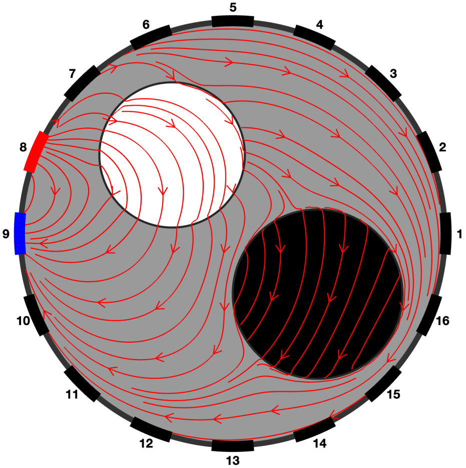

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

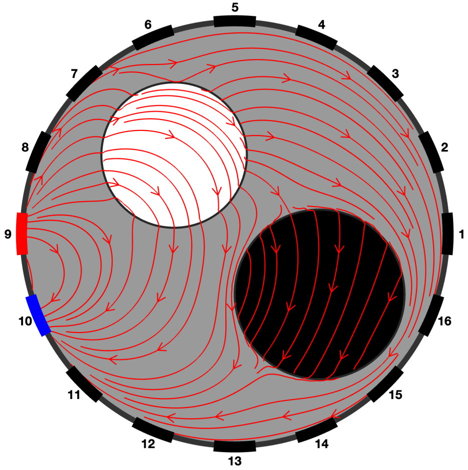

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

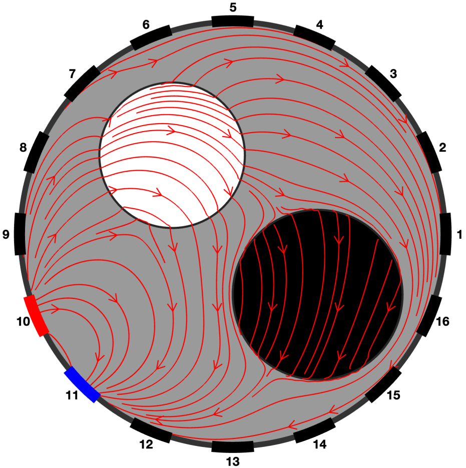

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

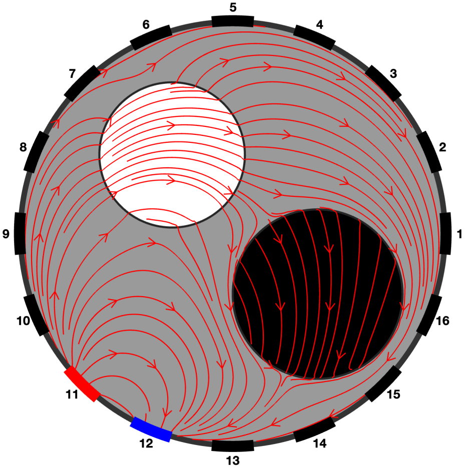

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

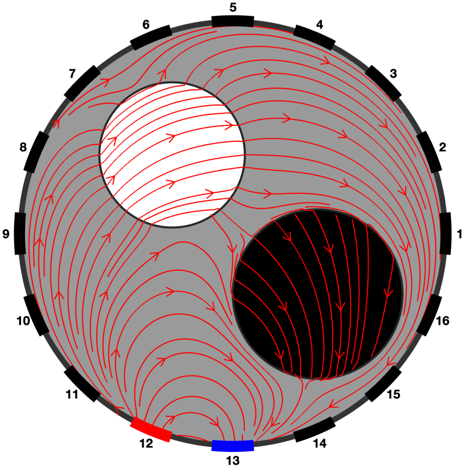

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

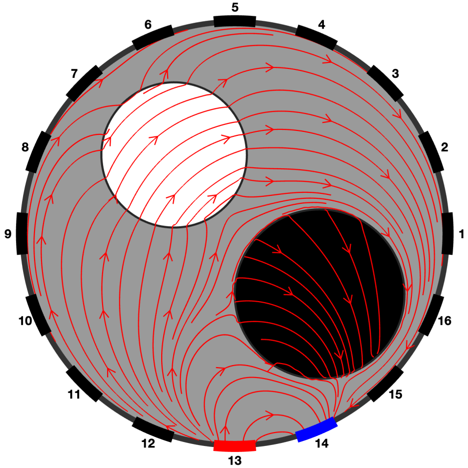

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

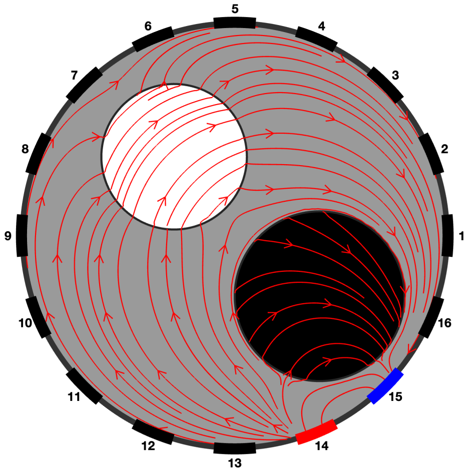

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

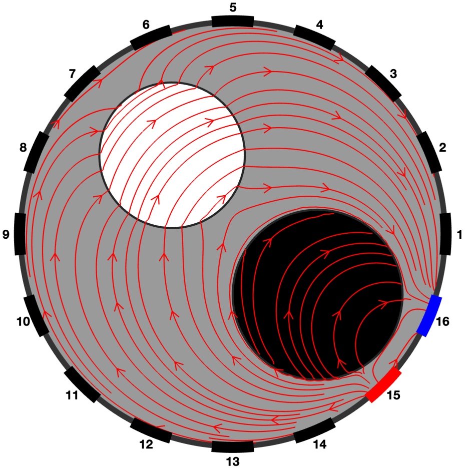

Rotating the Drive Pair

Suppose we take a series of measurements, rotating the drive pair around the domain.

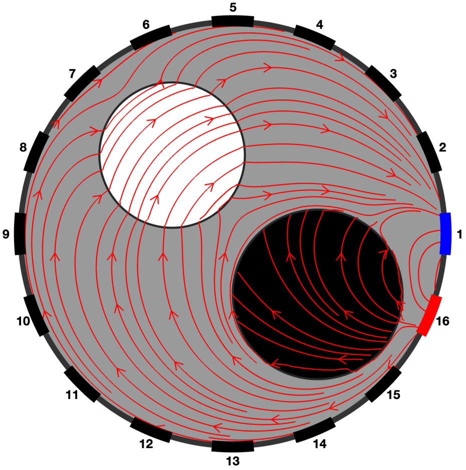

The Current and Voltage Data Matrices

Injection n produces one input-output pair:

$$\mathbf{I}^{(n)} \longmapsto \mathbf{V}^{(n)}.$$

Collecting the experiments column by column gives a current pattern matrix I and voltage matrix V:

$$\boxed{\mathsf{I} = [\mathbf{I}^{(1)}, \ldots, \mathbf{I}^{(N)}] \in \mathbb{R}^{L \times N}, \qquad \mathsf{V} = [\mathbf{V}^{(1)}, \ldots, \mathbf{V}^{(N)}] \in \mathbb{R}^{L \times N}.}$$

- Each column corresponds to one current injection (N of these).
- Each row corresponds to one electrode (L of these).

In our example, $L = N = 16$, so both matrices are 16×16.

Where Do These Measurements Go Next?

Together, the matrices I and V describe what the instrument applied and measured.

Different reconstruction methods use these data in different ways:

- Many methods work directly with the measured current and voltage data.
- CGO-based methods (e.g. D-bar) first convert them into a discrete matrix approximation of the ND or DN map.

We will develop that matrix construction later, when we introduce CGO-based reconstruction.