Intramural Forecasting of Cardiac Electrical Dynamics Using Data Assimilation

Matthew J. Hoffman and Elizabeth M. Cherry

School of Mathematical Sciences
Rochester Institute of Technology

Rochester, NY, USA

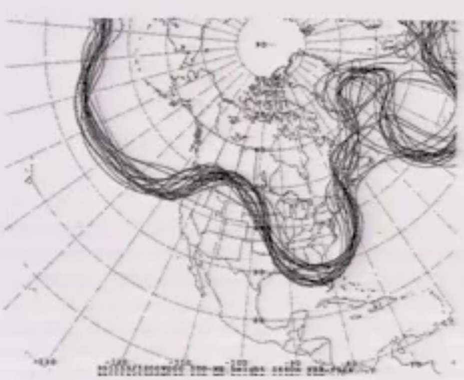

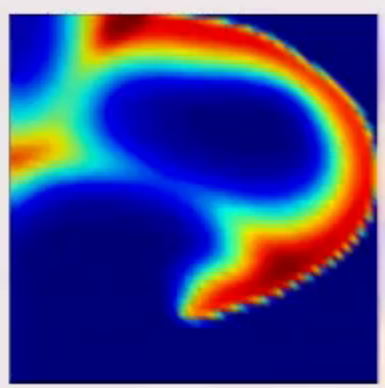

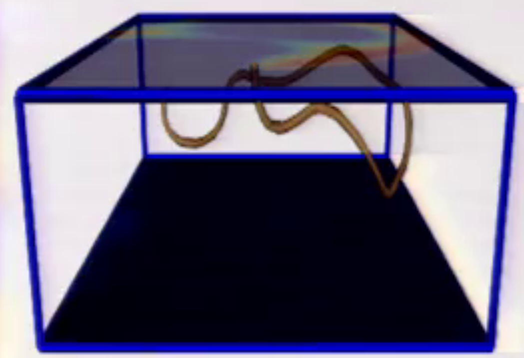

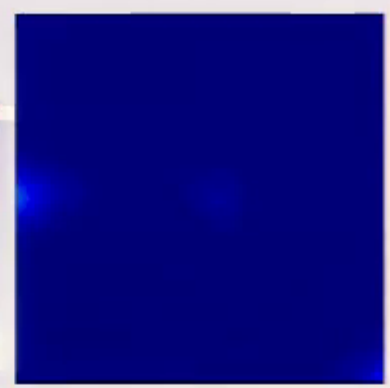

Motivation

- Arrhythmias are disorders of the regular propagation of electrical waves in the heart, often leading to reentrant spiral/scroll waves and spatiotemporal chaos.

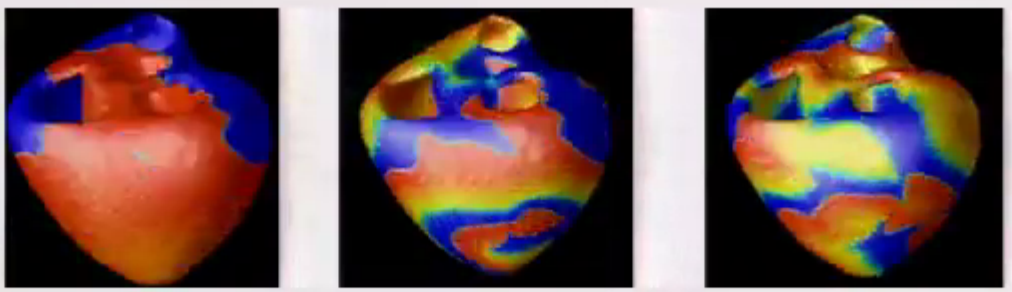

Motivation

- Arrhythmias are disorders of the regular propagation of electrical waves in the heart, often leading to reentrant spiral/scroll waves and spatiotemporal chaos.
- To study arrhythmias, we can record electrical activity from both outer and inner surfaces.
- Tissue thickness is ~1cm, but it is not always obvious how observations from the two surfaces are related.

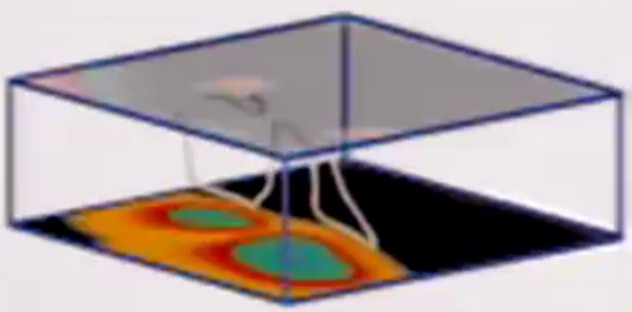

Endocardium

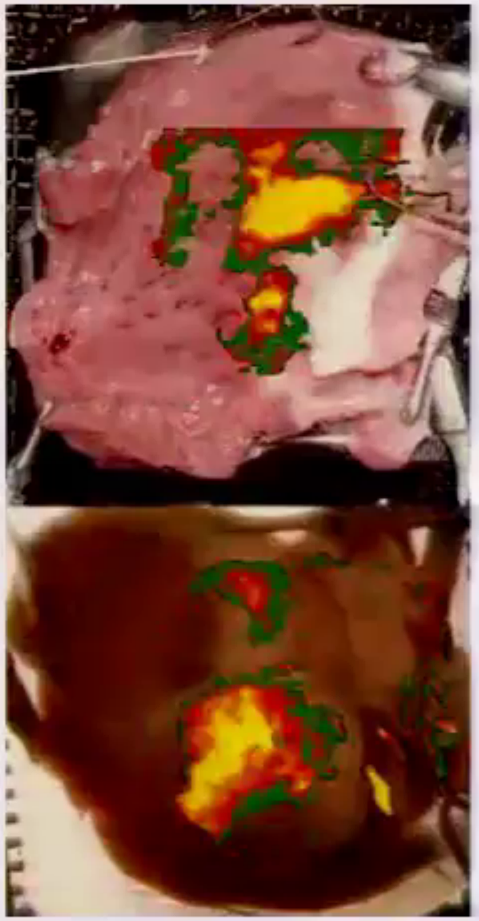

Epicardium

Fenton FH et al. 2008. New Journal of Physics 10, 125016.
Cherry EM and Fenton FH. 2013. Frontiers in Cardiac Electrophysiology 4, 71.

Motivation

- Arrhythmias are disorders of the regular propagation of electrical waves in the heart, often leading to reentrant spiral/scroll waves and spatiotemporal chaos.
- To study arrhythmias, we can record electrical activity from both outer and inner surfaces.
- Tissue thickness is ~1cm, but it is not always obvious how observations from the two surfaces are related.

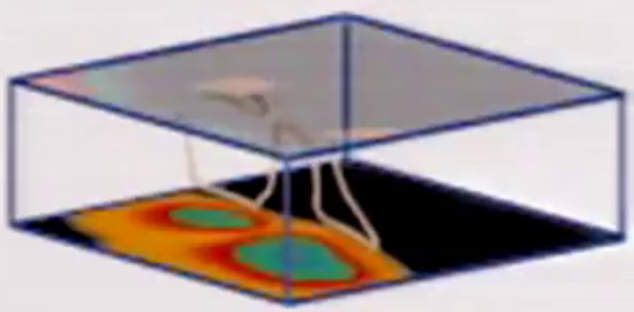

Endocardium

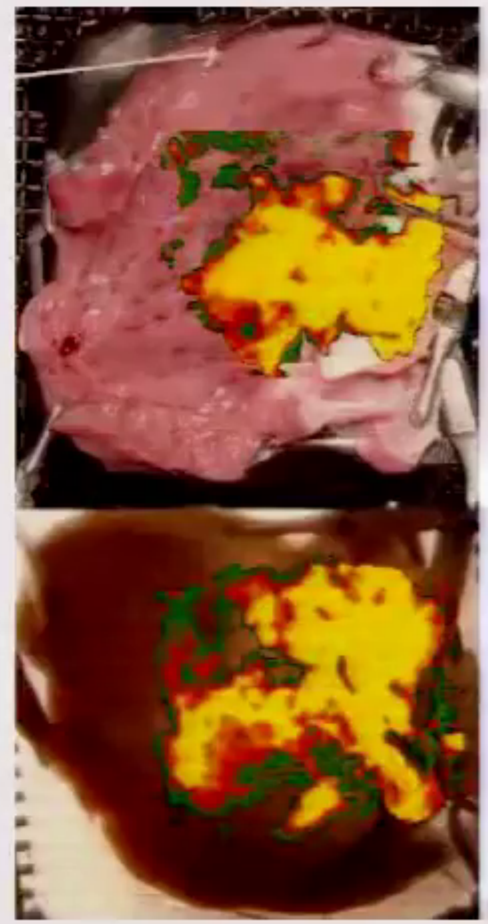

Epicardium

Fenton FH et al. 2008. New Journal of Physics 10, 125016.
Cherry EM and Fenton FH. 2013. Frontiers in Cardiac Electrophysiology 4, 71.

Motivation

- We want to reconstruct the 3-D propagation and breakup of electrical waves in cardiac tissue experiments.
- This 3-D time series may help us understand what is going on in the unobserved thickness of the muscle.
- Solving this problem requires both state estimation (voltages, concentrations, etc.) and the ability to forecast from a given state.
- "Data assimilation" is used for this purpose in the weather-forecasting community: observational data are combined with numerical model predictions to improve a forecast—or, here, a reconstruction.

Forecasting as an initial-value problem

- Mathematically, prediction requires both a model of the system,

$$\frac{\partial \mathbf{x}}{\partial t} = F(t, \mathbf{x}),$$

  and an estimate of the current system state, $\mathbf{x}(t_j)$.

- But there will be errors in both!
  - Model: errors in formulation (approximation) and numerical solution.
  - Initial state estimate: direct measurement generally is not possible.

State estimation

- Models and observations can both be used for state estimation (time series: series of state estimates).
- Models: "predict" the state by running the model.
  - Excellent spatial and temporal resolution.
  - Approximations of the true dynamics (quantitatively).
- Observations: interpolate from observed data.
  - Reflect the true system.
  - Spatially and temporally sparse.
  - Interpolating would ignore dynamics (we can use information from previous times to provide additional constraints).
- Data-assimilation approach: combine new observations with a model-derived state estimate, based on older observations.

Data assimilation cycle

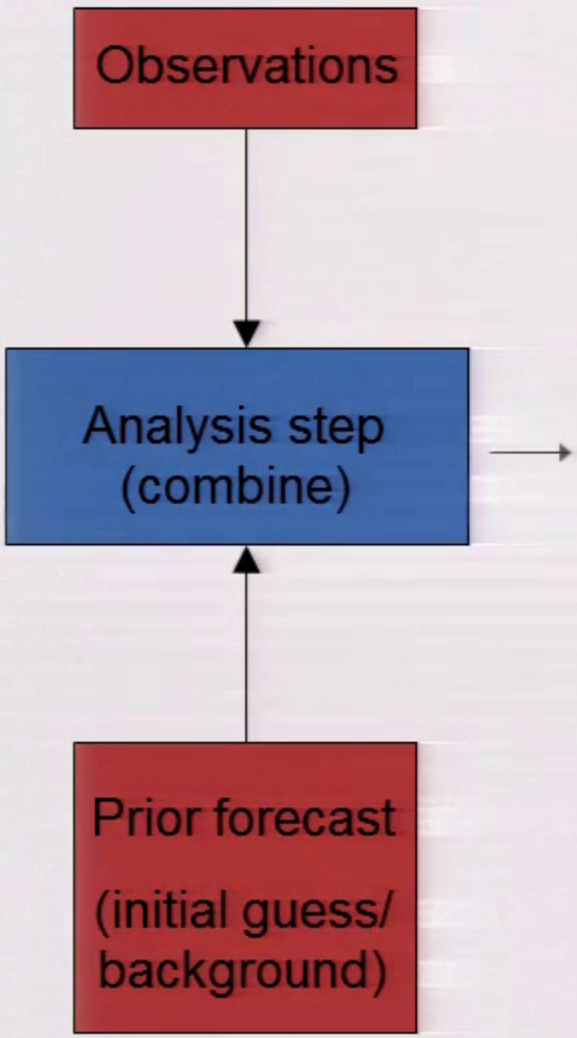

Data assimilation is an iterative process.

Data assimilation cycle

Data assimilation is an iterative process.

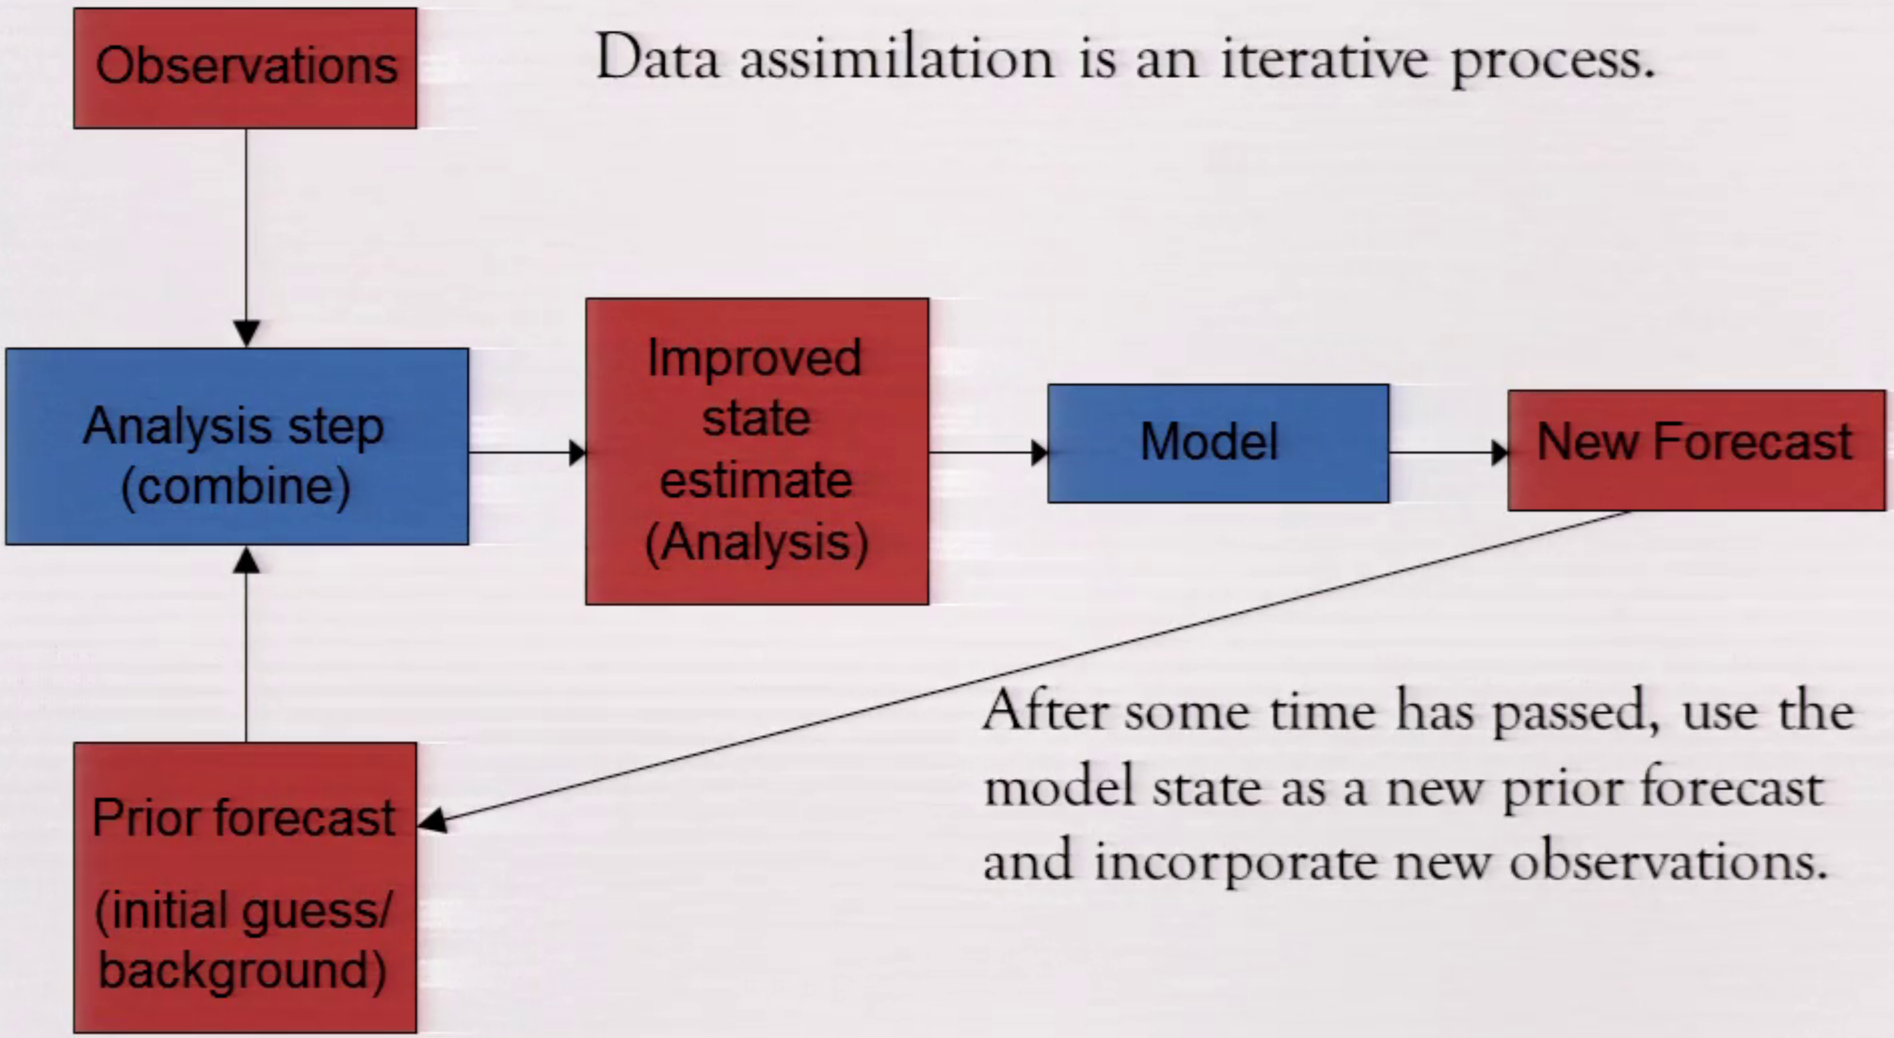

After some time has passed, use the model state as a new prior forecast and incorporate new observations.

Kalman Filter

- Most data-assimilation methods are based on the Kalman filter, which produces state estimates from noisy data.
- The Kalman filter (Kalman 1960, 1961) was designed for a linear model and observation operator.
- It assumes Gaussian observation errors: $\mathbf{y}_j^o = H\left(\mathbf{x}(t_j)\right) + \epsilon_j$.
- In the linear case (reasonable for many nonlinear models), the cost function becomes

$$J_{t_n}^o(\mathbf{x}) = [\mathbf{y}_n^o - \mathbf{H}_n\mathbf{x}]^T \mathbf{R}_n^{-1} [\mathbf{y}_n^o - \mathbf{H}_n\mathbf{x}] + [\mathbf{x} - \mathbf{x}_n^b]^T (\mathbf{P}_n^b)^{-1} [\mathbf{x} - \mathbf{x}_n^b] + c.$$

Ensemble Kalman Filter

- Model size makes computation with the background covariance prohibitively expensive (invert a huge matrix).
- Instead: use an ***ensemble*** of model states to characterize the background state and its covariance (common approach for representing uncertainty).

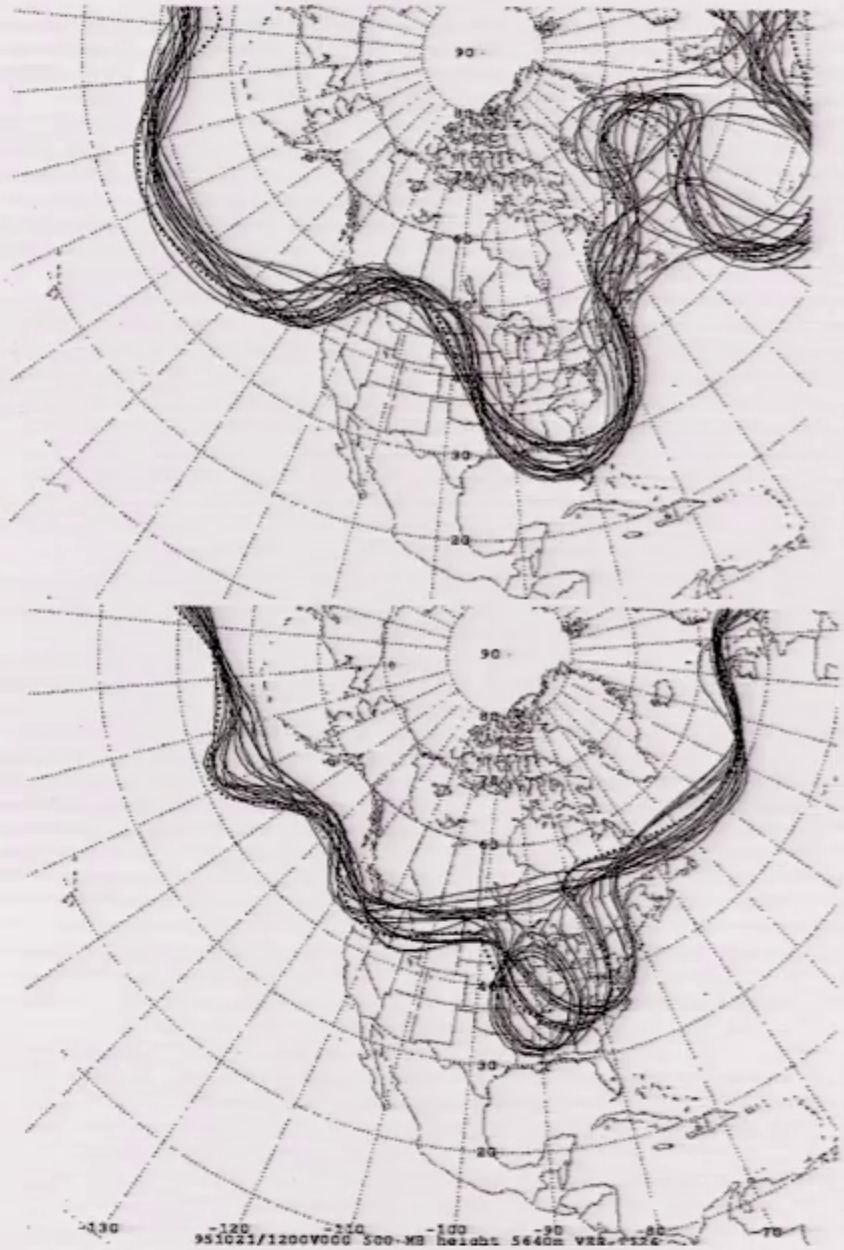

Ensemble Kalman Filter

- Start with a k-member ensemble at time $n-1$.

$$\left\{\mathbf{x}_{n-1}^{a(i)} : i = 1,2,\ldots,k\right\}$$

- Each ensemble member is propagated to the next time using the model (function M).

$$\left\{\mathbf{x}_{n}^{b(i)} = M\left(\mathbf{x}_{n-1}^{a(i)}\right)\right\}$$

Ensemble Kalman Filter

- Start with a *k*-member ensemble at time $n-1$.

$$\left\{\mathbf{x}_{n-1}^{a(i)} : i=1,2,\ldots,k\right\}$$

- Each ensemble member is propagated to the next time using the model (function *M*).

$$\left\{\mathbf{x}_{n}^{b(i)} = M\left(\mathbf{x}_{n-1}^{a(i)}\right)\right\}$$

- At that time, the "best guess" of the state from the model (background state estimate) is given by the ensemble sample mean.

$$\overline{\mathbf{x}_n} = \frac{1}{k}\sum_{i=1}^{k}\mathbf{x}_n^{b(i)}$$

- The background covariance is given by the ensemble sample covariance (how much variation in ensemble?).

$$\mathbf{P}^b = (k-1)^{-1}\sum_{i=1}^{k}\left(\mathbf{x}_n^{b(i)} - \overline{\mathbf{x}_n^b}\right)\left(\mathbf{x}_n^{b(i)} - \overline{\mathbf{x}_n^b}\right)^T = (k-1)^{-1}\mathbf{X}^b(\mathbf{X}^b)^T$$

Cardiac application

- We aim to reconstruct the 3-D times series of electrical wave propagation and breakup in cardiac tissue.
- Experimentally, observations of one variable (voltage) are available at the tissue surfaces (from cameras recording fluorescence signals).
- We first consider known simulated states to evaluate the potential of data assimilation in advance of testing using experimental data.
- A numerical prediction model is used.

Numerical model

- We use the 3-variable (u, v, w) Fenton-Karma model to update the voltage $u(t, \xi)$ by the sum of all transmembrane currents I_{ion} and diffusive coupling:

$$\frac{\partial u(t,\xi)}{\partial t} = \nabla \cdot D(\xi)\, \nabla u(t, \xi) - I_{ion}(t, \xi).$$

- $D(\xi)$ contains information about the arrangement of cells (rotational anisotropy in 3D).
- The system is solved numerically.

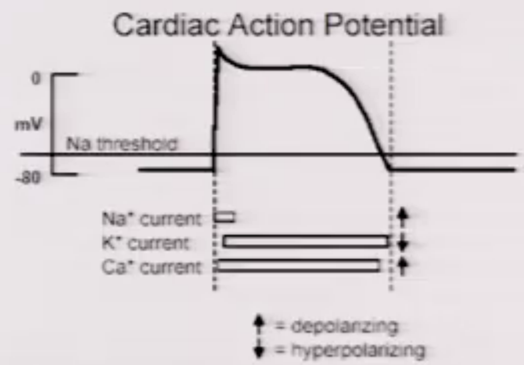

Fenton FH, Karma A. 1998. Chaos 8, 20-47.

Initial experiments

- ***Truth*** is given by the numerical simulation (a knowable state to allow for performance evaluation and testing).
- Synthetic ***observations*** are created by adding random Gaussian error to a subsampling (in space and time) of the truth.
- Using the same model both to generate truth and to evolve state estimates forward in time eliminates model error and puts the full focus on algorithm performance.
- We show initial results for both 1-D and 3-D.

1-D wave propagation

- In 1-D, the Fenton-Karma model is set up on a ring (14 cm, 0.025cm spacing) and the system is placed in a state with wavelength oscillations (discordant alternans).

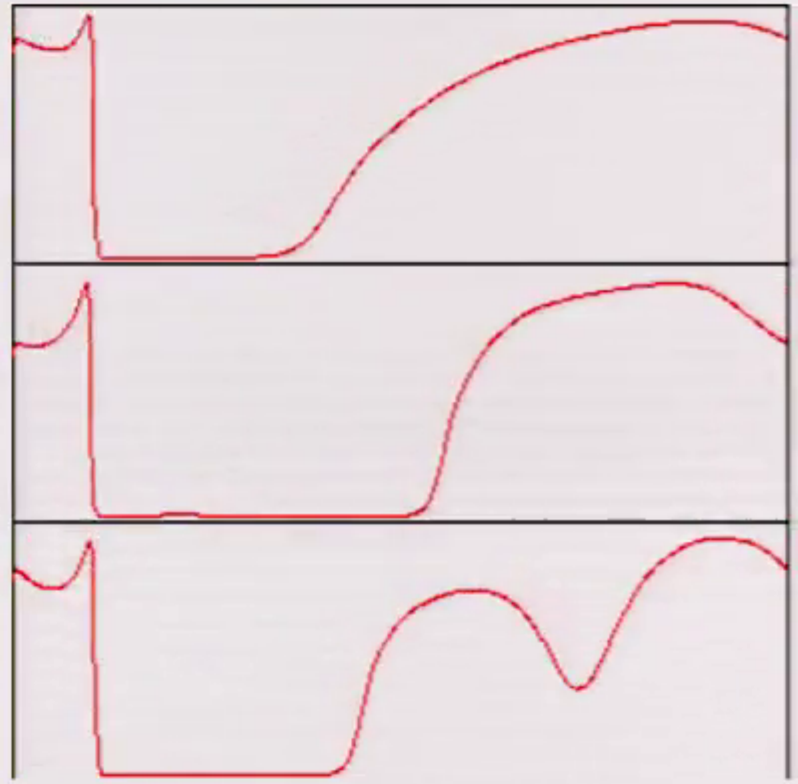

1-D wave propagation

- In 1-D, the Fenton-Karma model is set up on a ring (14 cm, 0.025cm spacing) and the system is placed in a state with wavelength oscillations (discordant alternans).

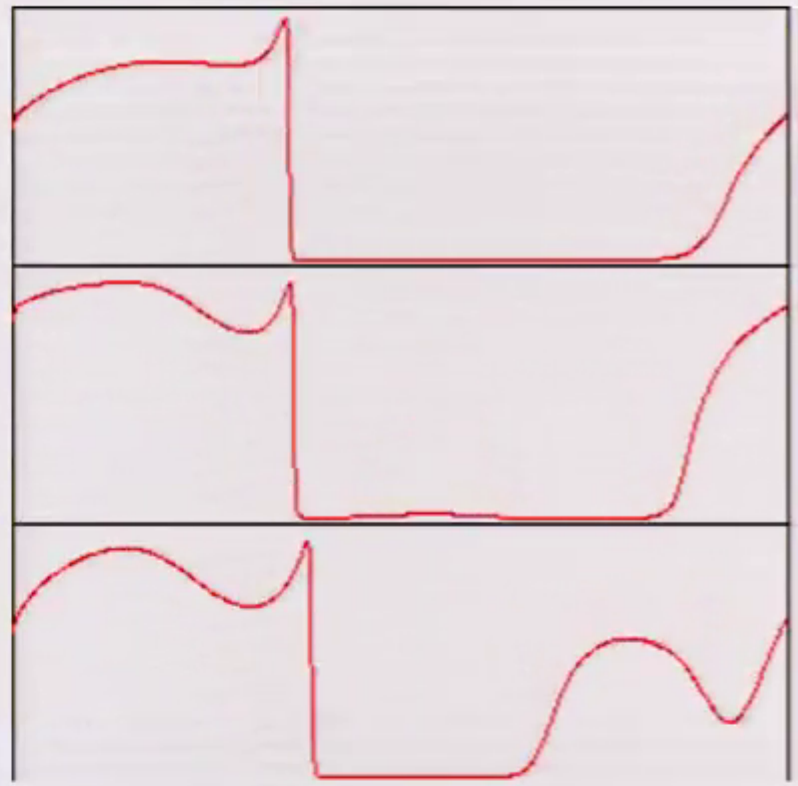

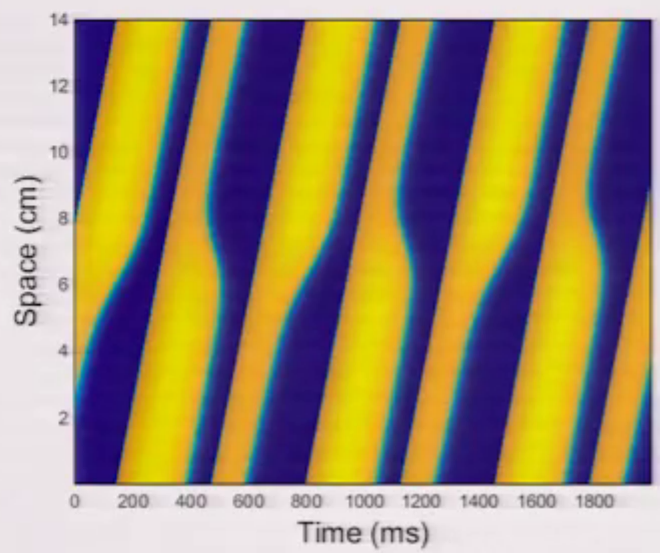

1-D wave propagation

- In 1-D, the Fenton-Karma model is set up on a ring (14 cm, 0.025cm spacing) and the system is placed in a state with wavelength oscillations (discordant alternans).

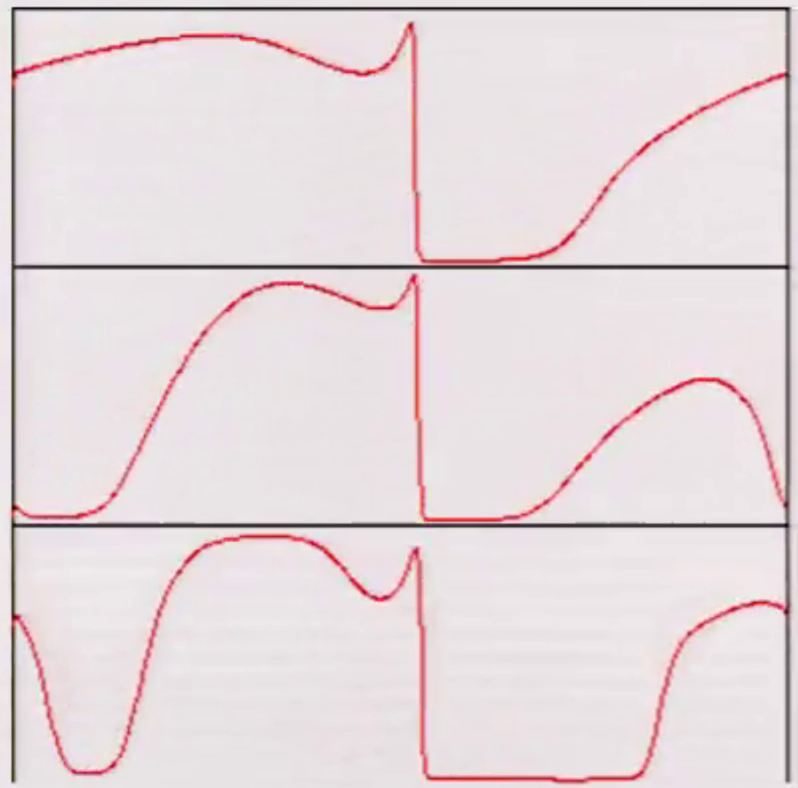

- The model is run to generate the "truth."
- Observations: random Gaussian error ($\sigma = 0.05$) added to the voltages of this "truth" every 5 ms, grid spacing 0.075cm).
- Resolution is comparable to (or worse than) typical cameras.

1-D test run

- Assimilation: every 5 ms.
- 20-member ensemble (blue), initialized to states from 40ms prior to first assimilation plus random Gaussian error ($\sigma = 0.05$).
- Note that the ensemble mean (black) has a different front structure from any of the ensemble members (blue).
- u and v are corrected by u observations, but not w (variable localization).
- Multiplicative inflation is used (increase covariance artificially) with $\rho = 1.2$.
- Localization is used with σ=0.05 cm (observations are used within a radius of $2\sqrt{10/3}\,\sigma$).

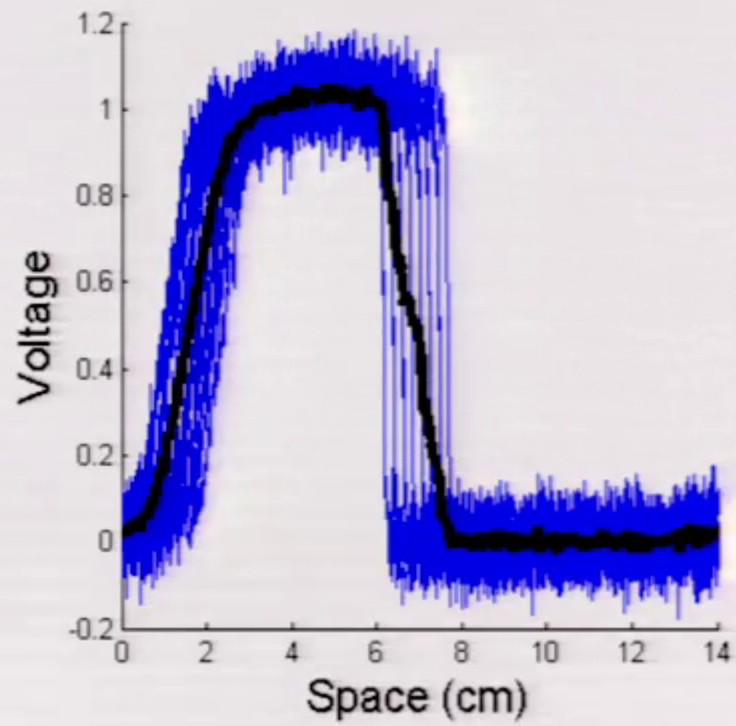

1-D test run

- After the first assimilation, the analysis (observations incorporated) is a much better fit to the **truth** and observations than the background (left figure).
- The analysis quickly converges to the **truth** and the RMS error remains low throughout the 2-second simulation (right figure).

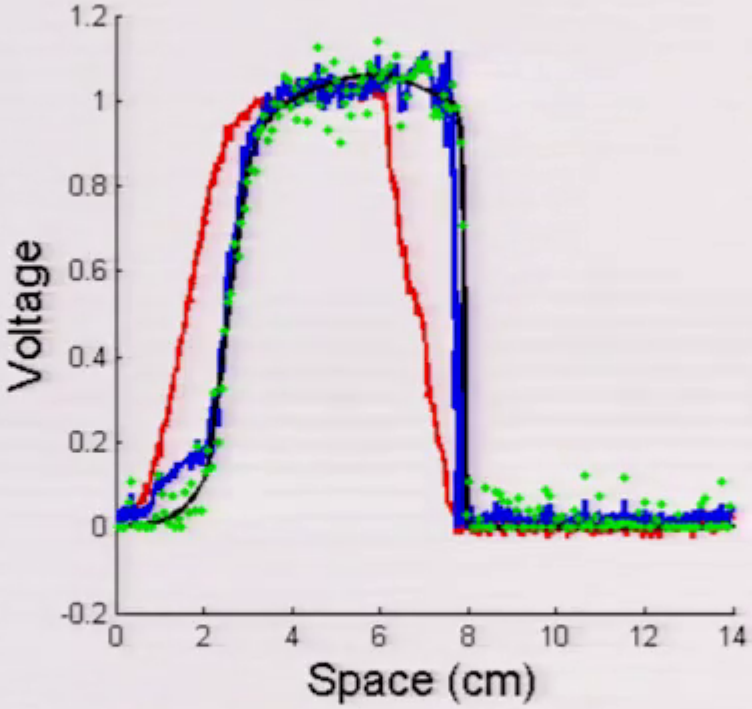

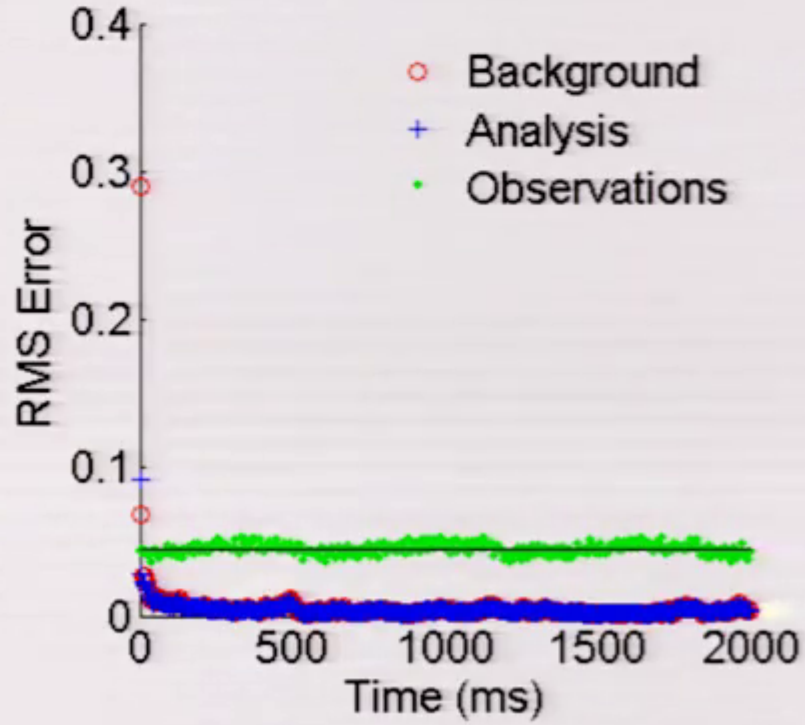

Response to different initial conditions

- When the initial ensemble is a poor estimate of the truth, the assimilation initially can fail.
- Example: 20-member ensemble, initialized to states from 1000ms prior to first assimilation.
- The wave now dies in the forecast and is unable to recover, even with observations of the truth.

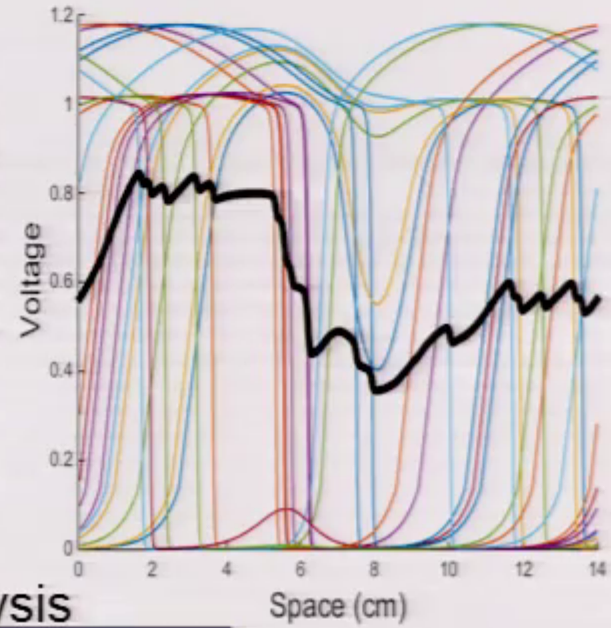

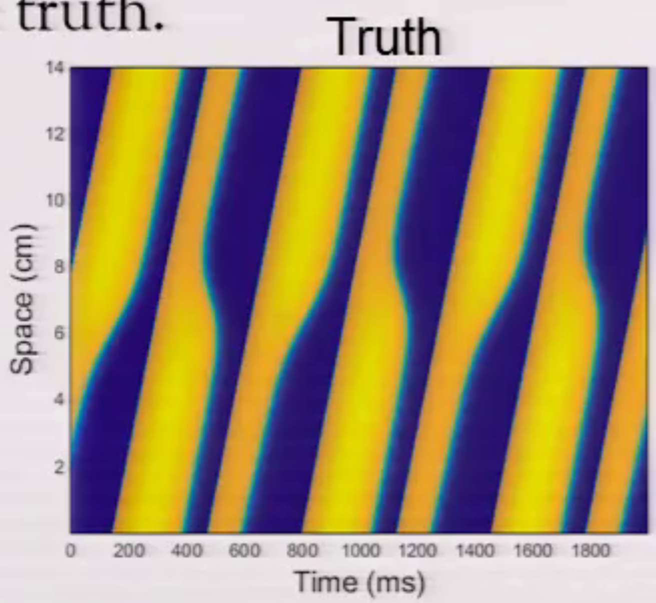

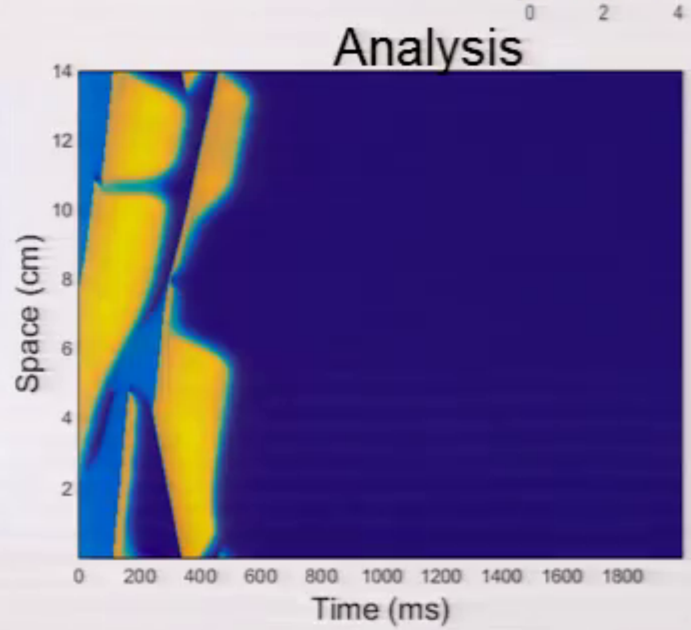

Response to different initial conditions

- When the initial ensemble is a poor estimate of the truth, the assimilation initially can fail.
- Example: 20-member ensemble, initialized to states from 1000ms prior to first assimilation.
- The wave now dies in the forecast and is unable to recover, even with observations of the truth.

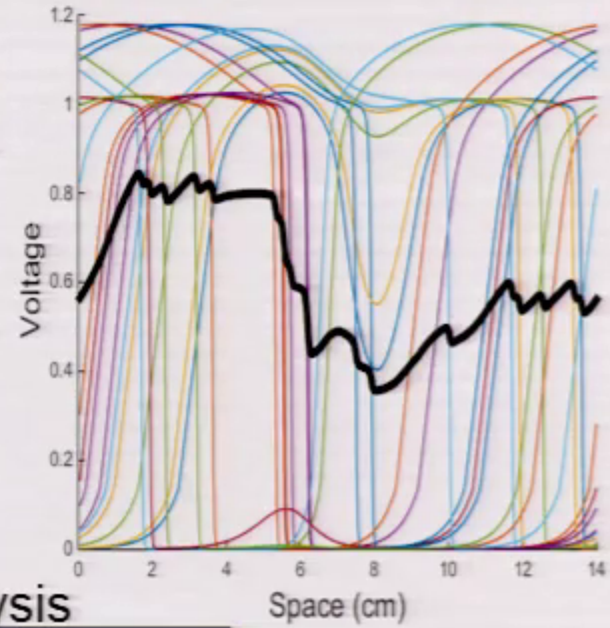

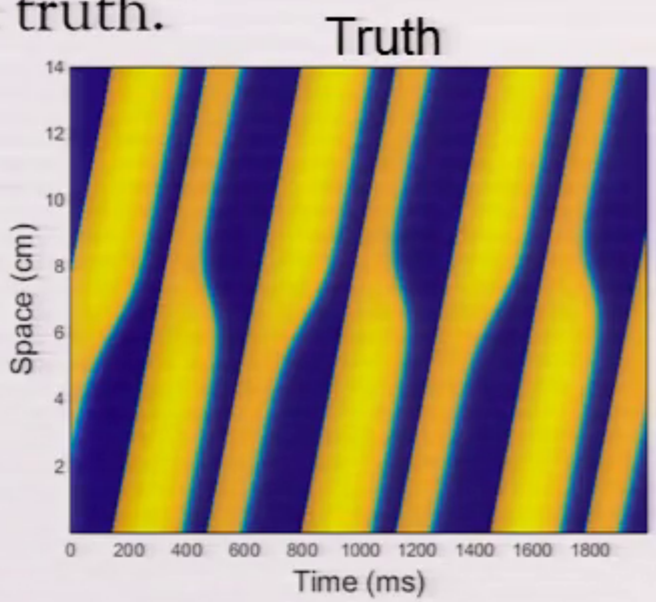

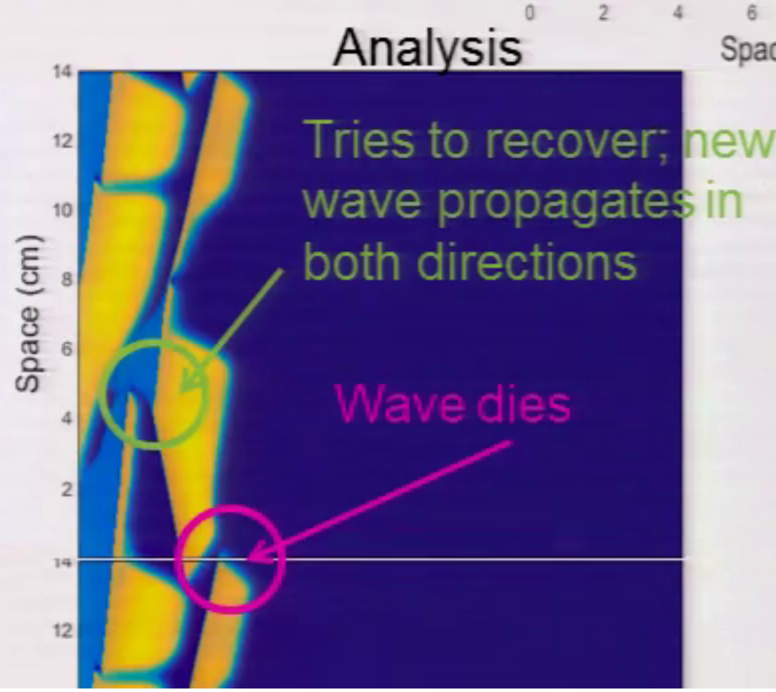

Ensemble collapse

- The system is unable to correct itself because the ensemble members become very similar, leading to overconfidence in the background and an inability to respond to observations.
- Multiplicative inflation tries to manage this, but cannot add new dimensions to the background.

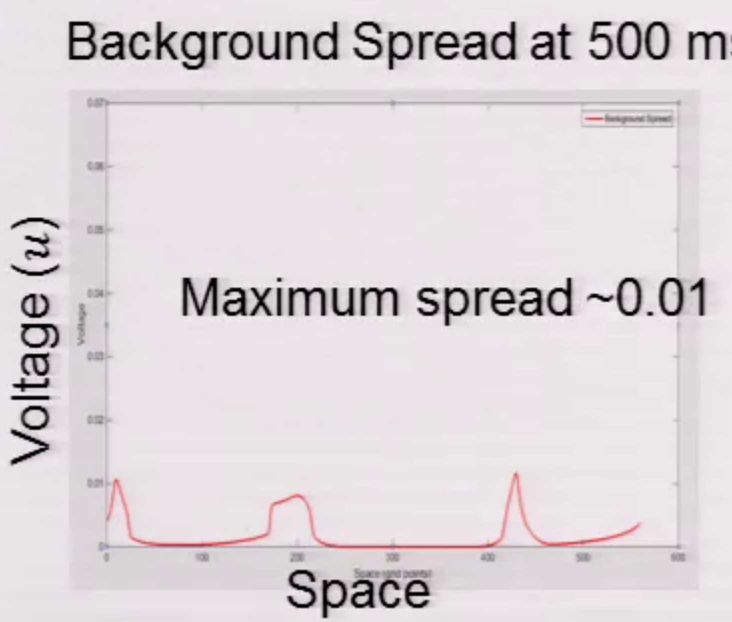

Ensemble collapse

- The system is unable to correct itself because the ensemble members become very similar, leading to overconfidence in the background and an inability to respond to observations.
- Multiplicative inflation tries to manage this, but cannot add new dimensions to the background.
- Instead, new vectors can be *added* to the ensemble to not only increase spread, but also change the space spanned by the ensemble.

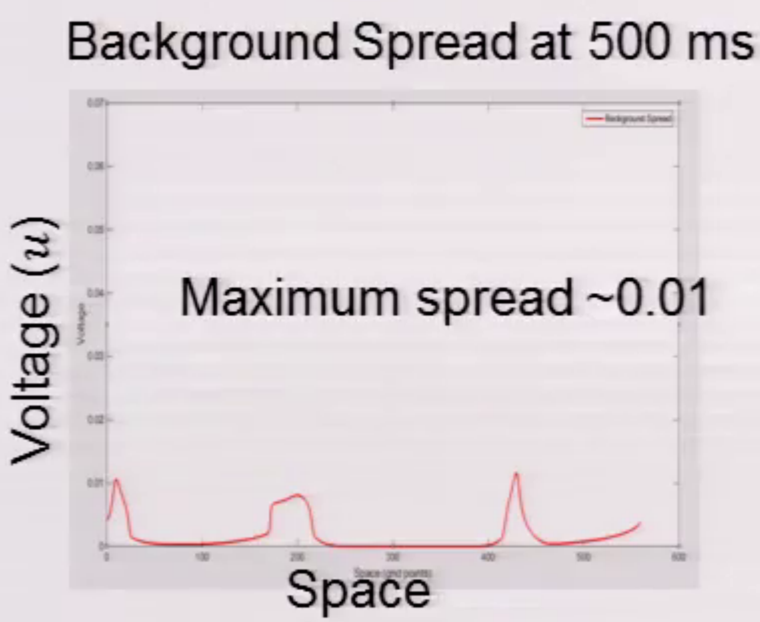

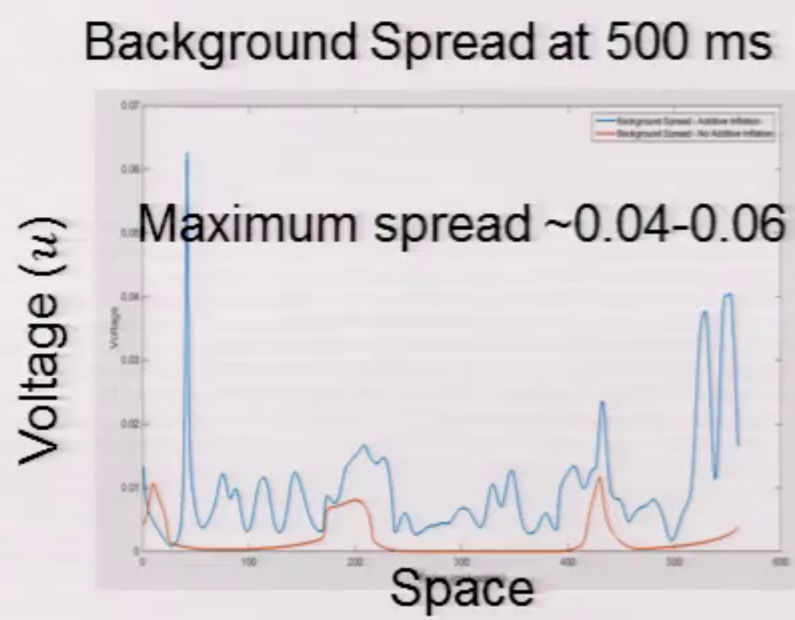

Additive inflation

- The additive inflation allows the algorithm to recover after the initial wave dies.
- Now after about 500 ms, the system syncs with the truth and remains close to it.
- The RMS error stays below the observation error after the initial 500ms.

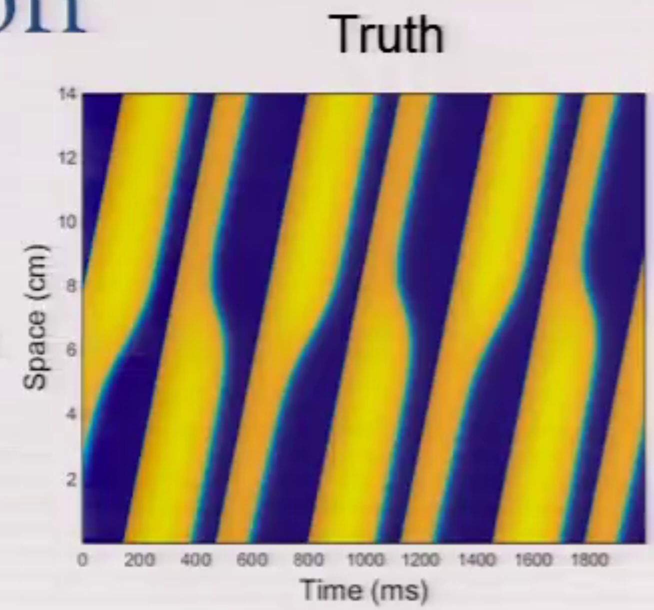

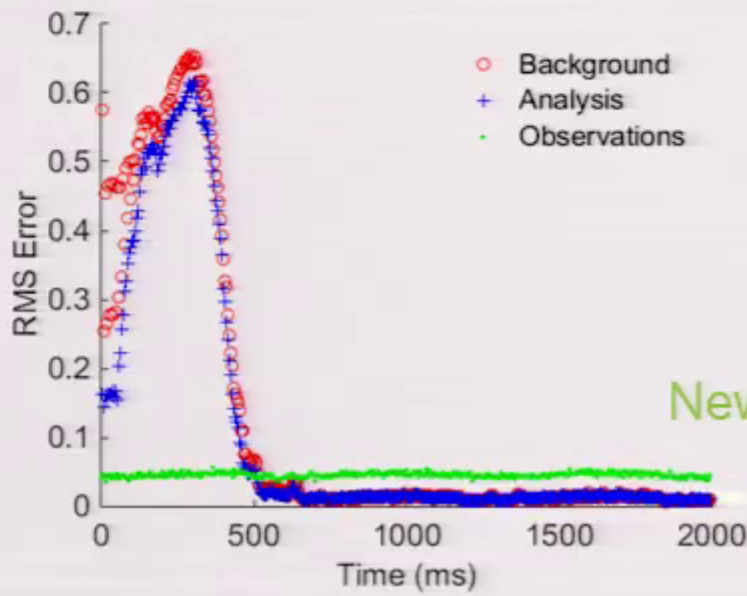

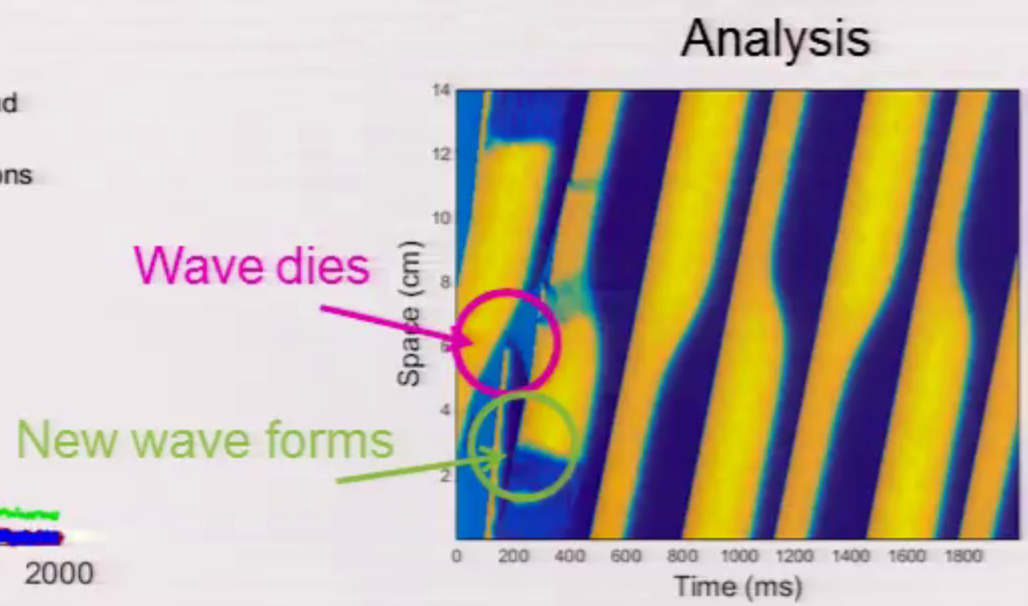

Unobserved variables

- Although only u observations are used, both unobserved fields converge to the truth within 500 ms.
- This is encouraging for real experiments where v and w cannot be observed.
- Both u and v are corrected; these corrections indirectly correct w.

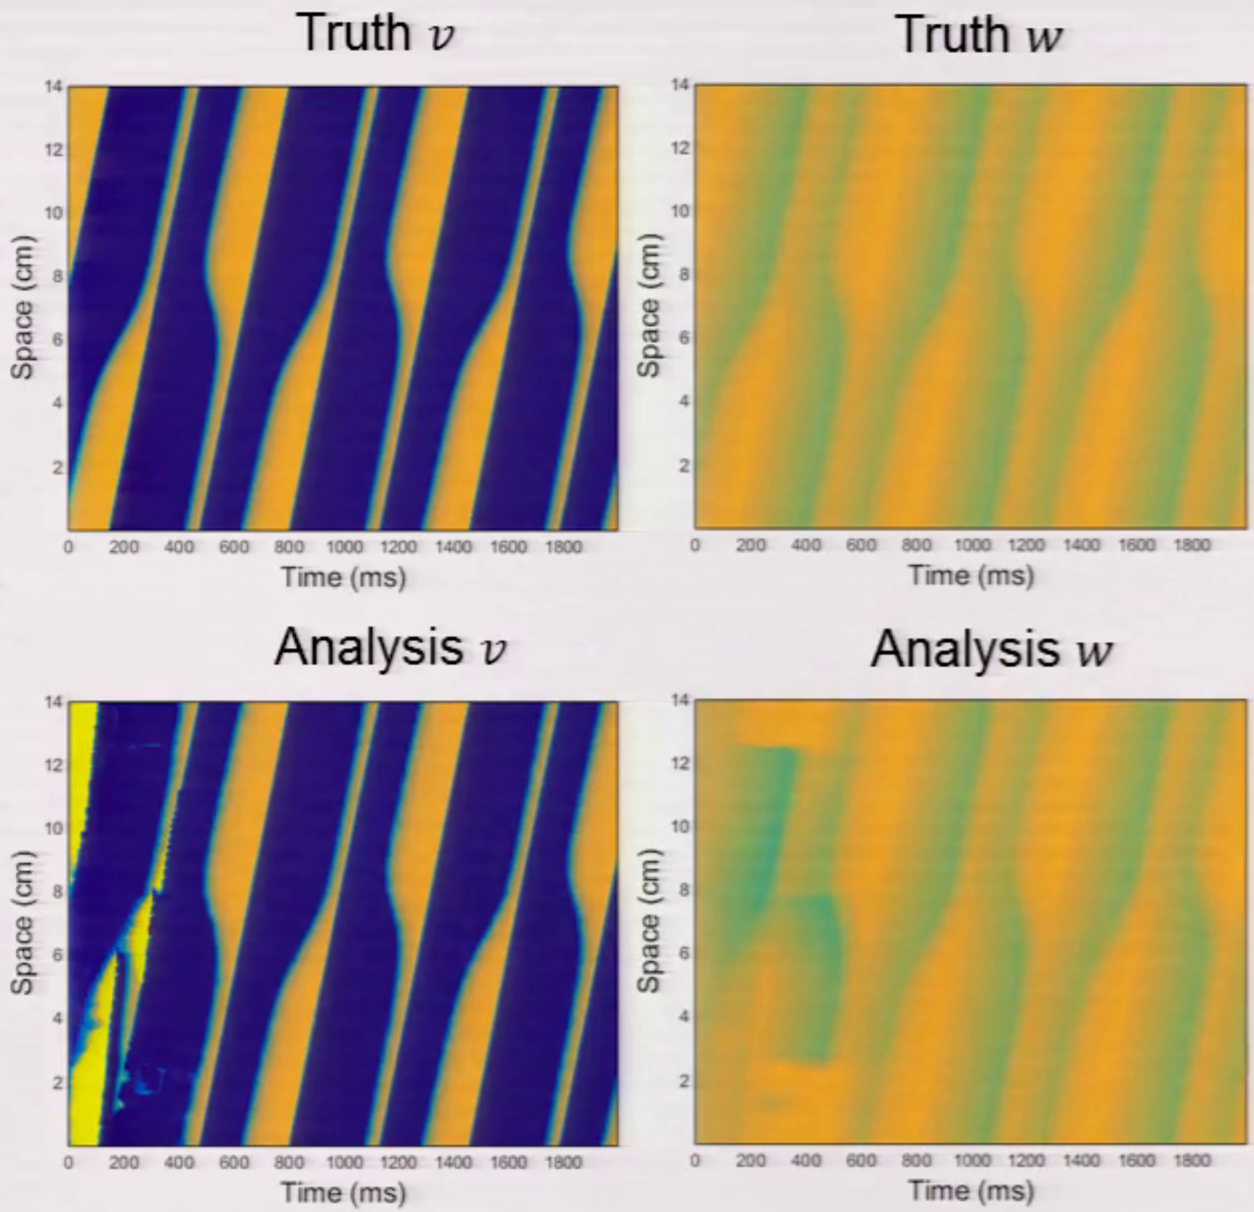

3-D setup

- We use a stack of spiral waves to start a scroll wave.
- 20 ensemble members are used based on the 1-D results.

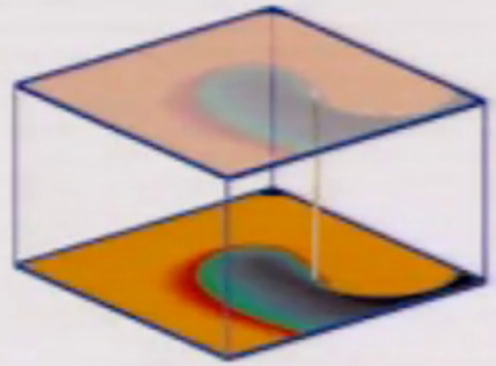

- Localization is used (σ=3 grid point, so an 18-grid-point radius of influence).
- Assimilation is performed every 5ms (as in 1D) and using every 3 grid points (0.06 cm observation grid spacing).
- The initial ensemble is generated using the previous 20 model states 5ms apart from the spinup.
- Multiplicative inflation factor is 1.1.

3-D LETKF results

- As expected, the initial guess is poor due to the initialization.
- The initial analysis significantly improves the voltage estimate and recovers most of the scroll wave.

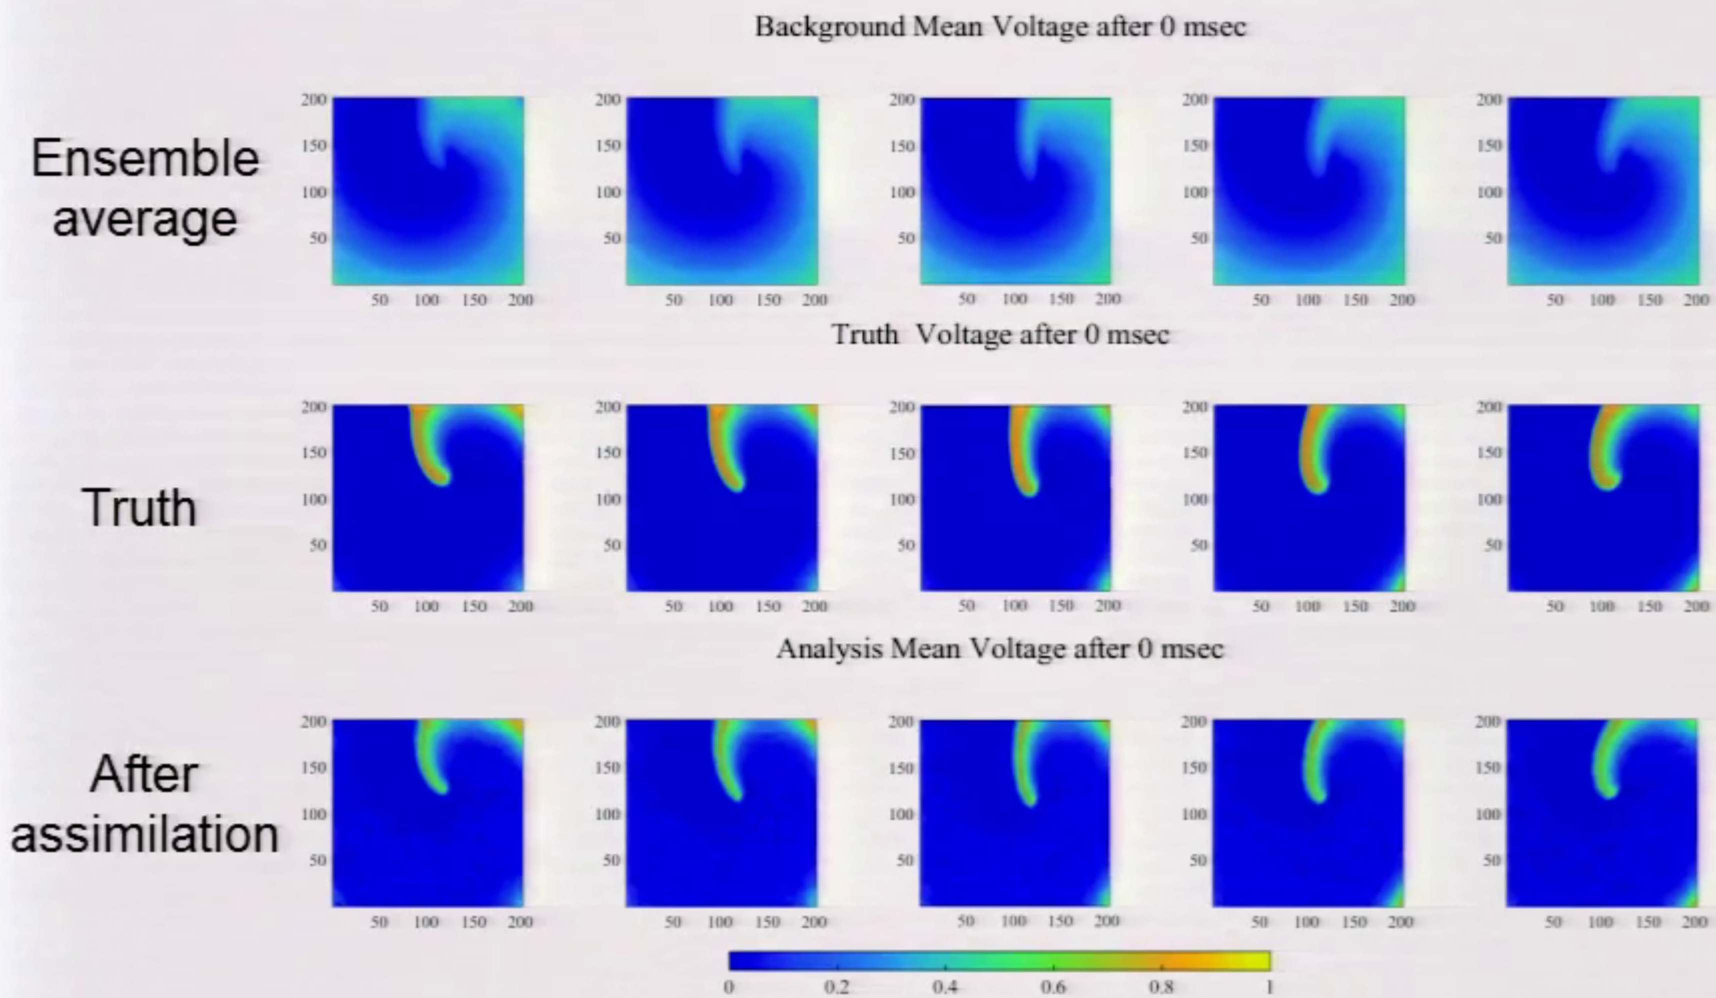

3-D LETKF results

- After several assimilation cycles the system converges to the truth.
- The convergence in again seen in all variables even though only u is observed.
- Here u, v, and w are corrected.

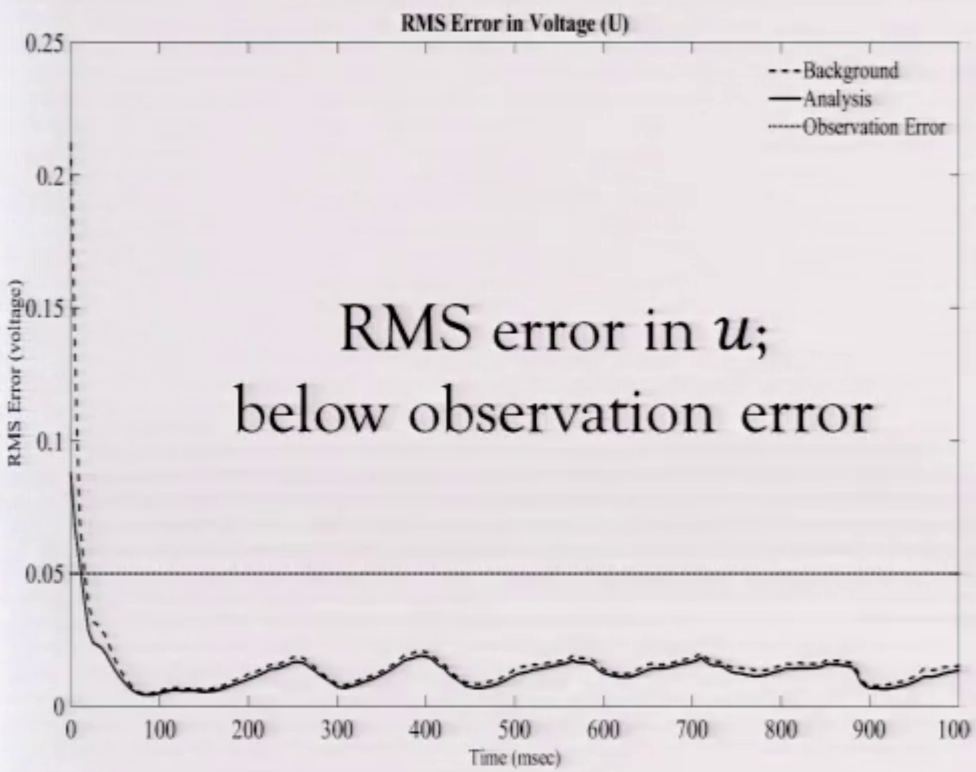

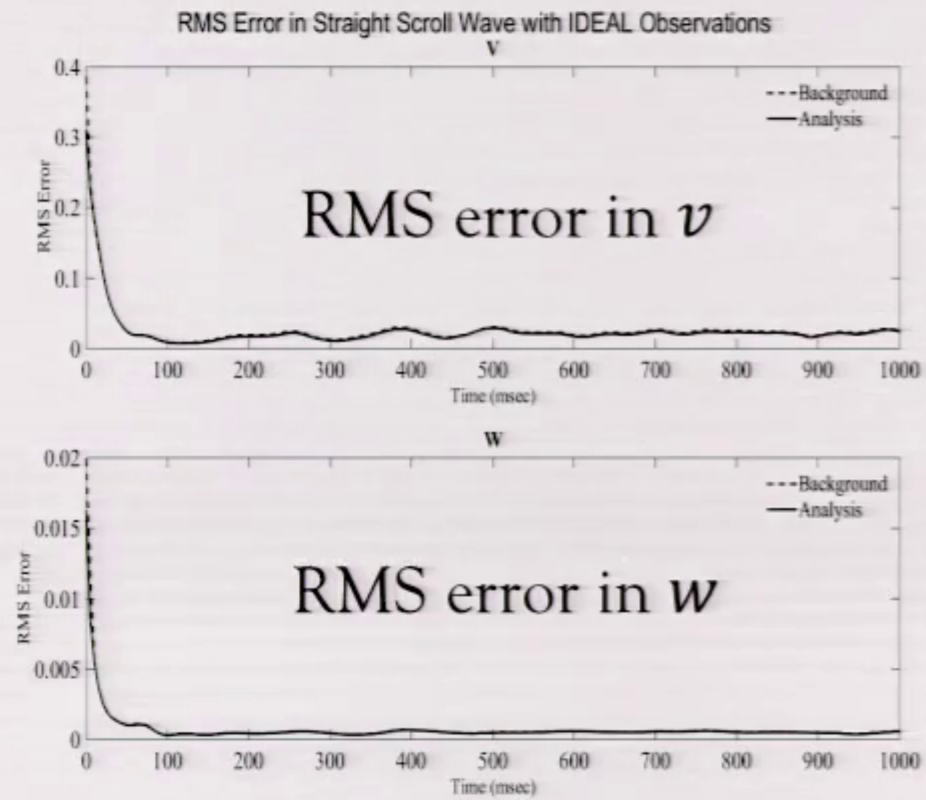

3-D LETKF results

- The largest errors in the analysis come from two areas:
  1. Lower voltage in the center of the wave (yellow).
  2. Smoothing of the sharp wave front (blue).
- The analysis does correct both, but not enough—yet!

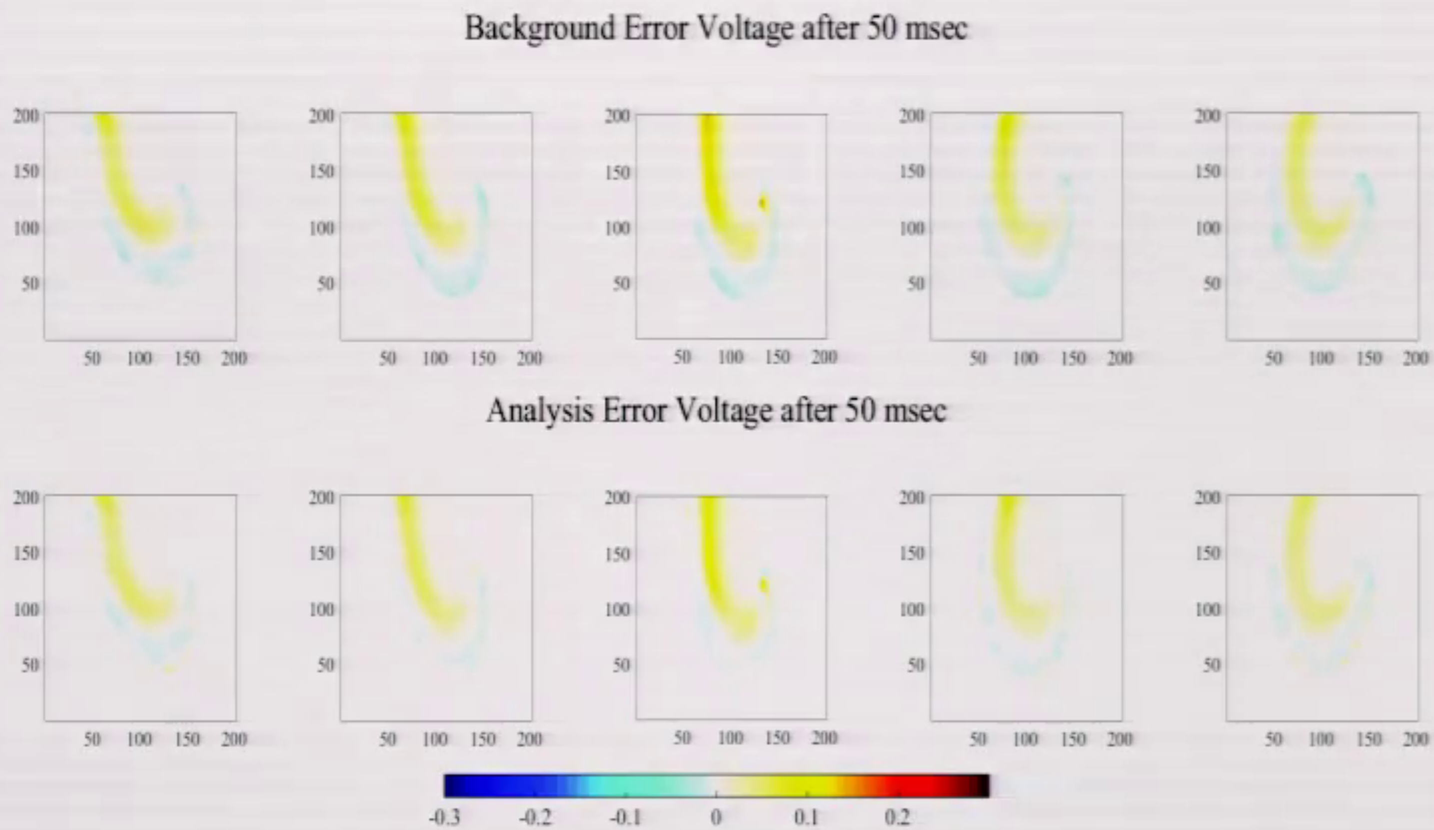

Summary

- Data assimilation shows promise as a means of reconstructing 3D time series in cardiac applications.
- Findings thus far:
  - Fairly low-dimensional space (20 ensemble members).
  - Additive, but not multiplicative, inflation confers ability to recover from very poor initial guess.
  - Corrections based on observations in one variable successfully correct other variables.
  - Some sensitivity to initialization.

Ongoing and future work

- Analyze how the initial states chosen affect robustness.
- Study the effects of model error.
- Consider more complicated dynamical states.
- Use more realistic 3-D observation distributions to see how far into the interior information can be propagated reliably.
- Use the algorithm to estimate model parameters—we have begun testing this capability.
- Investigate ways to initialize and simulate real tissue.