Iterative Learning - Gone Wild

LLRF 2022, Brugg, CH
Shane Koscielniak; 2022 Oct 07

TRIUMF-Elinac
30 MeV, 10 mA c.w. capable
1.3 GHz SRF
electron linear accelerator

EGUN: electron gun
ELBT: low energy transport
EINJ: injector cryomodule

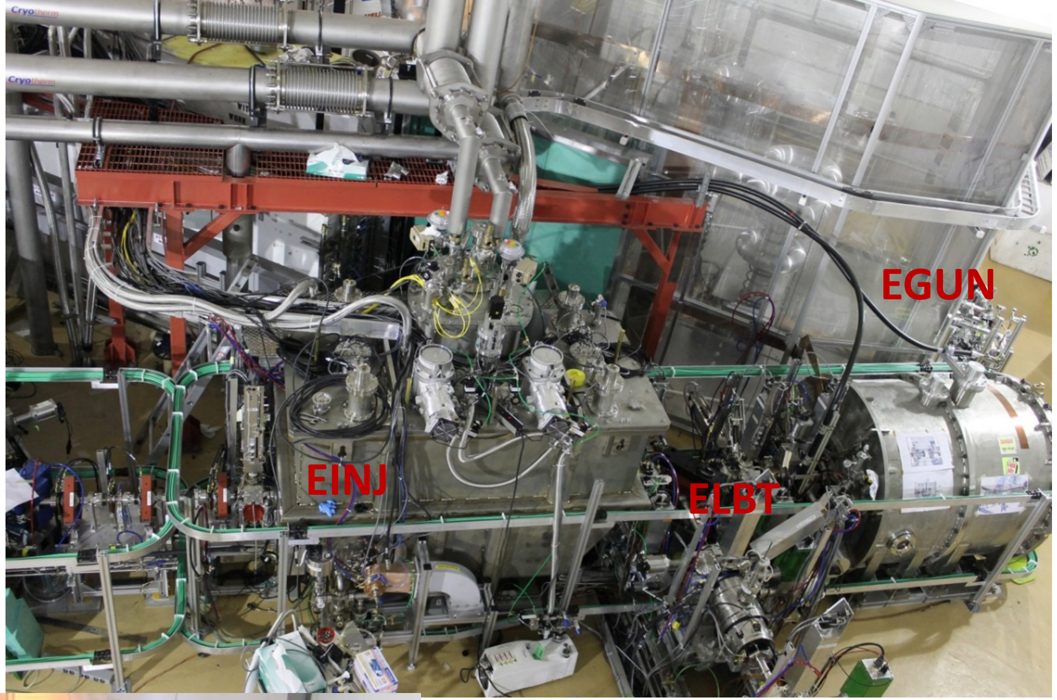

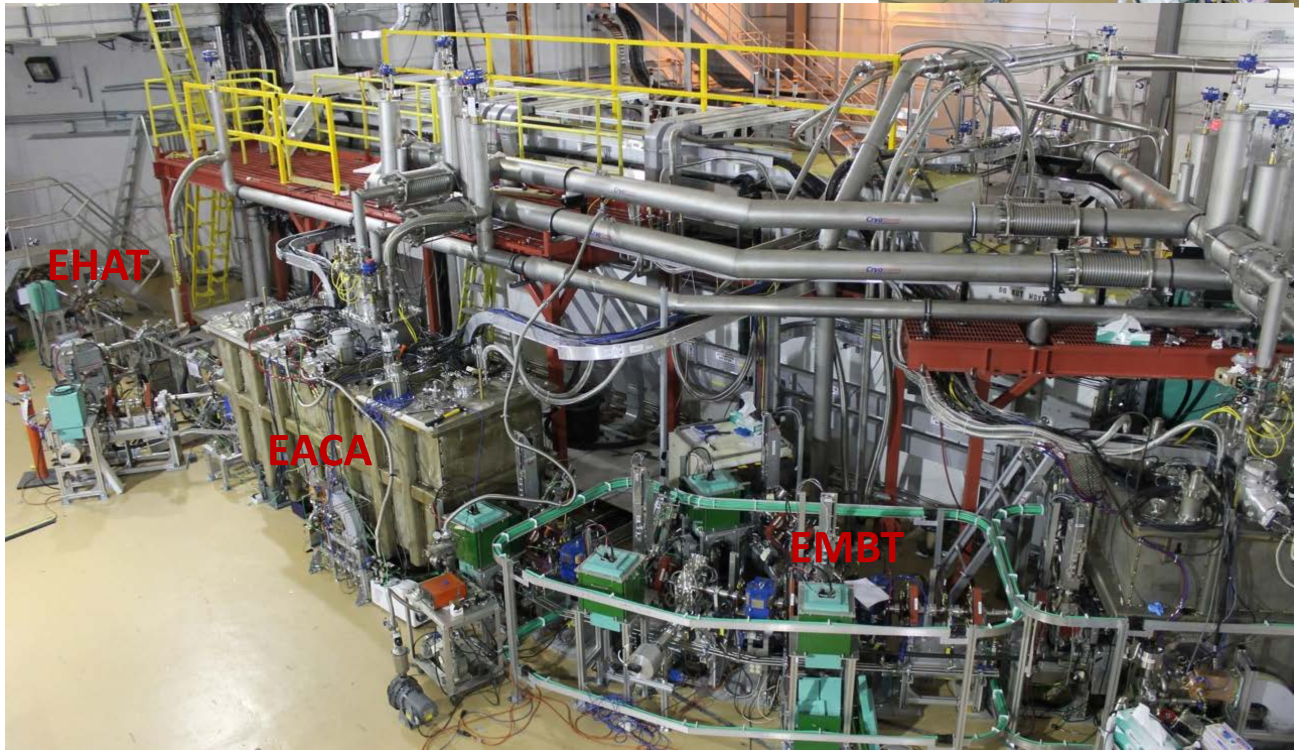

EMBT: medium energy transport
EACA: accelerator cryomodule EHAT: accelerated beam transport

- TRIUMF E-linac is nominally a c.w. accelerator
- But must operate at much lower power for commissioning and beam development
- So e-linac is pulsed with a variety of repetition rate and pulse length, leading to significant transient beam-loading of the SRF cavities
- Feed-forward compensation (FFC) of beam-loading transients is essential.
- A different FFC[t] function is needed for every combination of repetition rate and pulse length
- In 2013, LLRF group proposed (TRI-DN-13-23) to generate the many FFC[t] time functions by Iterative Learning*.
- Based on numerical simulation, they chose a 4-term noncausal ILC.
  - Q=I and L=I+⅓(↑+↑↑+ ↑↑↑)
  - But numerical cases do not guarantee stability or convergence, because they cannot span the entire space of initial conditions
- So present author, embarked on a complete analysis, culminating in the discovery in 2016 of wave solutions and prediction of unstable learning.

*See, for example, the review:
IEEE Control Systems Magazine, Vol.26, No.3, 2006, pp. 96-114

Principle of Iterative Learning Control (ILC)

- System (or plant P) with its own internal (stable) response.
- Place an ILC wrapper around P and iterate wrapped system from one trial to next.
- During the trial, the plant is a free system with a driven input.
- At the end of an individual trial, the input is updated based on results from the ILC wrapper.
- "Learning" = Matrix-map iteration of input to output vector.
    - Nested within the matrix is the internal response of the plant during a trial, which is different each trial. In time domain the iterations are non-linear.
- Hopefully the plant and its inputs & outputs and the map all "settle down" so that the wrapped system converges.

    - The E-linac LLRF system is a digital system, and well suited to implementing ILC
        - The RF waveform is demodulated and sampled
        - And all control signals are sampled
        - All samples can be recorded and held as super-vectors.

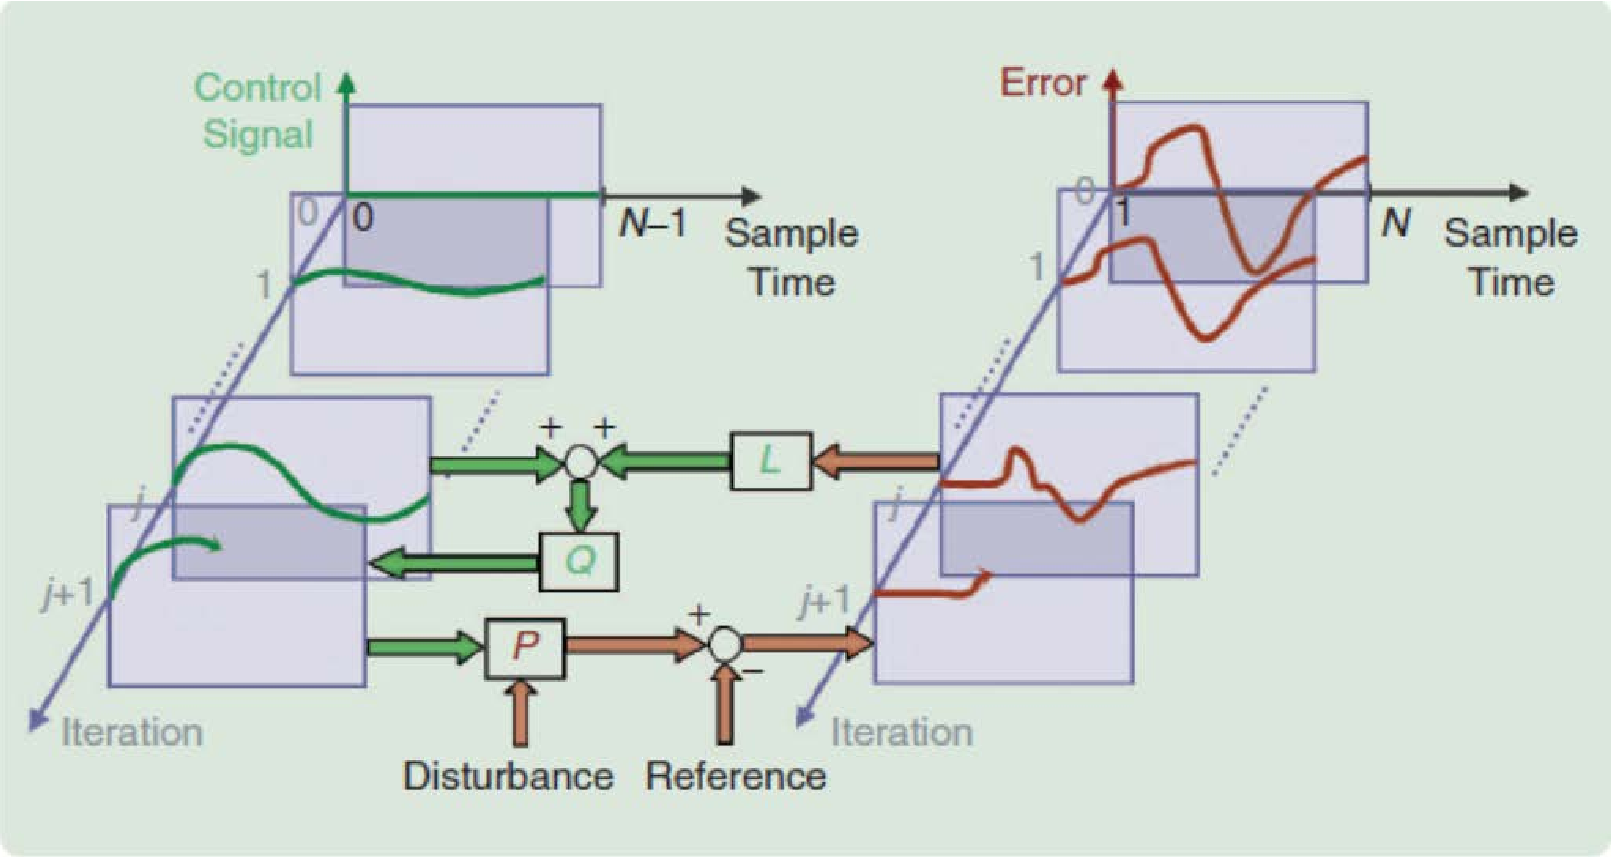

ILC Two dimensional space:

Within a trial, index k

From one iteration to the next, index j or n

Symbols & Equations

Vectors (internal index k)

u = input

y = output

d = disturbance (repetitive)

$\mathbf{y_d}$ = desired output (repetitive)

$\mathbf{e_j=(y_d\text{-}y)_j}$ = error

Operators

P = "the plant", i.e. the system

Q = filter

L = learning function

During the trial: $\mathbf{y_j=P_u u_j + P_d d}$

From one trial to the next, "**u**" learns from the previous trial (or trials).

$\mathbf{u_{j+1}= Q[u_j+Le_j]}$ where $\mathbf{e_j=(y_d\text{-}y)_j}$

Elimination of **e** leads to iterative maps for **u** and **y**.

These maps are function of the internal gains (K) and time constants (τ) of **P**, and the iteration gains (ν) of **L**; or the set {K,τ;ν} for short.

Principle of Iterative Learning Control

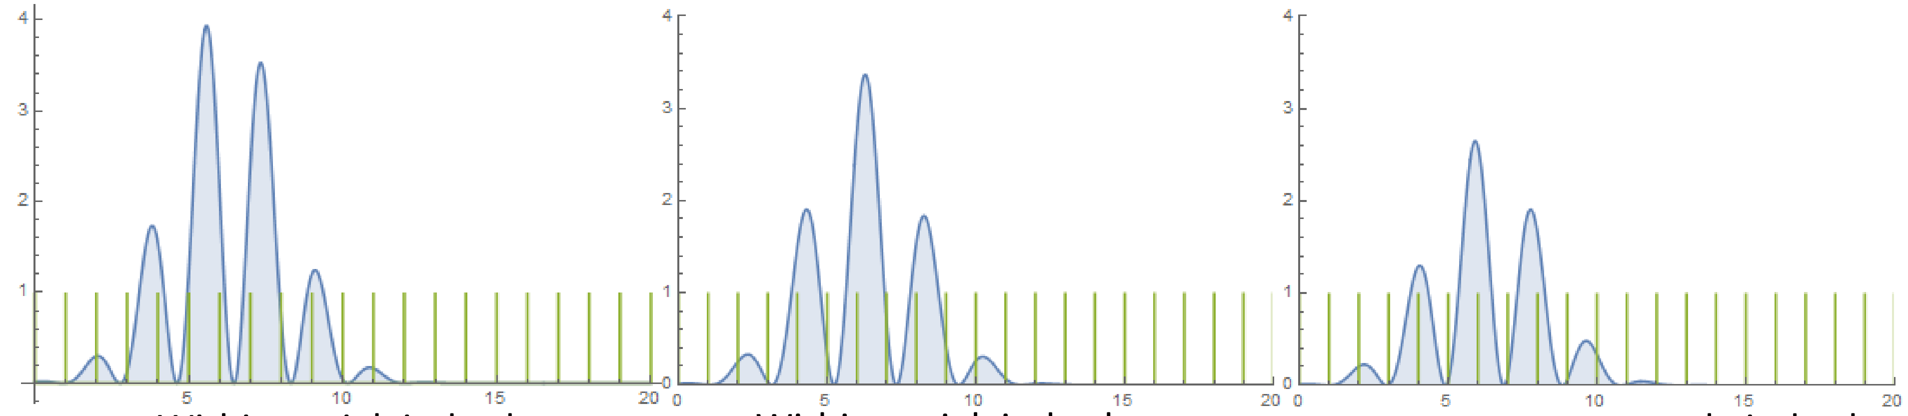

Within a trial, index k — Within a trial, index k — sample index k

$\mathbf{y_{j-1}=P_u u_{j-1} + P_d d}$ — $\mathbf{y_j=P_u u_j + P_d d}$ — $\mathbf{y_{j+1}=P_u u_{j+1} + P_d d}$

(j-1)T time (j-1)T<t<jT — jT time jT<t<(j+1)T — (j+1)T time (j+1)T<t<(j+2)T

iteration index j-1 — iteration index j — iteration index j+1

$\mathbf{u_j= Q[u_{j-1}+Le_{j-1}]}$

$\mathbf{u_{j+1}= Q[u_j+Le_j]}$

Converges to what?

Mapping: operator M such that x_{n+1}=M:x_n where x_n may be scalar or vector.
If M=M[n], mapping is non-linear.

Mappings have fixed points which may be stable or unstable, and are not necessarily zeros of x. Within the basin of attraction, a map will (eventually) converge on a stable fixed point.

Convergence is not a synonym for "stable".

"Stability" answers the question: "what happens when x is perturbed from its fixed point?".

The perturbations are infinitesimal, the system is locally linear, and the response is a decaying oscillation if "stable".

Iterations* converge on the fixed points:

$$[I-Q(I-LP_u)]u = QL(y_d-P_d d)$$

$$[I+P_u(I-Q)^{-1}QL\,]y = P_d d + P_u(I-Q)^{-1}QLy_d$$

$$[I+P_u(I-Q)^{-1}QL\,]e = y_d-P_d d$$

* If they converge

The case of **Q=I**
Fixed points simplify:

e=0, No residual error

$$y=y_d$$

$$u=P_u^{-1}(y_d-P_d d)$$

Iteration about the fixed points governed by:

$\tilde{\mathbf{u}}_{\mathbf{j+1}}$**=Q[I – LP]**$\tilde{\mathbf{u}}_{\mathbf{j}}$

$\tilde{\mathbf{y}}_{\mathbf{j+1}}$**=PQ[I – LP]P**$^{-1}$ $\tilde{\mathbf{y}}_{\mathbf{j}}$ **= [PQP**$^{-1}$ **– PQL]**$\tilde{\mathbf{y}}_{\mathbf{j}}$

Matrices of causal operators commute.

If the transfer functions **Q,P** are causal, then

$\tilde{\mathbf{u}}_{\mathbf{j+1}}$**= Q[I – LP]**$\tilde{\mathbf{u}}_{\mathbf{j}}$ and $\tilde{\mathbf{y}}_{\mathbf{j+1}}$ **= Q[I – PL]**$\tilde{\mathbf{y}}_{\mathbf{j}}$**.**

$\tilde{u}_j = u_j - u_{fp}$ & $\tilde{y}_j = y_j - y_{fp}$ & $\tilde{e}_j = e_j - e_{fp}$.

are "difference quantities" relative to the fixed points.

Properties of ILC maps

Usually it is stated (or assumed) that ILC maps generate sequences $\mathbf{x}_n$ that exist in a Banach space X; that is $\mathbf{x}_n$ is d-Cauchy convergent and the fixed point $\mathbf{x}$ is in X.

d-Cauchy convergence: for every positive real number r > 0 there exists some index N such that the distance $d(x_m,x_n)<r$ whenever m and n are greater than N

$||\mathbf{x}_n - \mathbf{x}||=0$ as n->infinity where $\mathbf{x}$ is in X.

The d-Cauchy property implies "eventually convergent", or asymptotically convergent.

Noncausal Learning Functions, and Exponential Convergence

- Sufficiently close to its fixed point, iterants x_n are exponential series (λ^n)
- $\tilde{\mathbf{y}}_{j+1} = \lambda \tilde{\mathbf{y}}_j$ implies the eigenvalue equation $\mathbf{Q[I - PL]}\tilde{\mathbf{y}}_j = \lambda \mathbf{I}\, \tilde{\mathbf{y}}_j$
- The limit of Convergence (of the iterations) means finding conditions {K,τ;ν} such that all eigenvalues (λ) of **A=Q[I – PL]** have modulus < 1.
- Of course, you may search for values {K,τ;ν} such that $|\lambda| < \gamma < 1$
  - $|\lambda(\mathbf{A})| < \gamma < 1 \rightarrow$ permissible {K,τ;ν}
- Coincidentally, the condition for exponential convergence is also the condition for Asymptotic Convergence of an ILC algorithm.
- Note: the within-trial motion has been decoupled from the iterations
- Away from the fixed point, other sequences x_n may satisfy the mapping.
- People have found them experimentally: the so-called learning transients
  - They persist long after the geometric sequences (λ^n) should have decayed
  - They are consistent with d-Cauchy convergence
  - They occur both for causal and noncausal learning functions
  - The "bad" initially diverge to large values before eventually converging

Causal Learning Functions, and Learning Transients

Suppose you look for solutions to the iteration equations in the form $\tilde{\mathbf{y}}_{j+n} = \lambda^n \tilde{\mathbf{y}}_j$ where $\mathbf{Q}[\mathbf{I} - \mathbf{PL}]\tilde{\mathbf{y}}_j = \lambda \mathbf{I}\, \tilde{\mathbf{y}}_j$

When **Q, P, L** are causal, **A**=**Q**[**I**-**PL**] takes on the Toeplitz form leading to a single repeated eigenvalue. In such case, iterants do not follow simple geometric series.

Solutions of the iteration equations may be obtained as the product of exponentials and power series.

$$\boxed{\mathrm{sol}[k,n] == \lambda^n \sum_{m=1}^{k} n^{k-m} \lambda^{-k+m} a[k,m]\mathrm{sol}[m,0]}$$

So, if you were expecting a geometric series, **all solutions are learning transients**

The within-trial behavior is coupled to the single eigenvalue.
Each element, sol[k], of the solution vector is coupled to preceding elements sol[k-m] & m<k-1

Asymptotic Convergence: As $n \rightarrow \infty$ each element sol[k,n] converges to zero provided the single eigenvalue $|\lambda|<1$.

The Element k does not start to converge until n>k; with $1 \le k \le N$ where N is matrix size

Simple Example: **Q=L=I**

A=(1-λ)

$$P=\begin{pmatrix} A & 0 & 0 \\ AB & A & 0 \\ AB^2 & AB & A \end{pmatrix}$$

$$\mathbf{x}[\mathrm{n}] = \begin{pmatrix} \mathrm{x1}[n] \\ \mathrm{x2}[n] \\ \mathrm{x3}[n] \end{pmatrix}$$

$$\{\mathrm{x1}[n] \to \lambda^n \mathrm{x1}[0]\}$$

$$\{\mathrm{x2}[n] \to -ABn\lambda^{-1+n}\mathrm{x1}[0] + \lambda^n \mathrm{x2}[0]\}$$

$$\left\{\mathrm{x3}[n] \to \frac{1}{2}\lambda^{-2+n}(A^2B^2(-1+n)n\mathrm{x1}[0] - 2ABn\lambda(B\mathrm{x1}[0] + \mathrm{x2}[0]) + 2\lambda^2\mathrm{x3}[0])\right\}$$

Each eigen-solution has a corresponding eigen-vector: {1,00}, {0,1,0}, {0,0,1}

Simple Example: **L=I+↓**

$$L=\begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix} \quad \mathbf{x}[\mathrm{n}]=\begin{pmatrix} \mathrm{x1}[n] \\ \mathrm{x2}[n] \\ \mathrm{x3}[n] \\ \mathrm{x4}[n] \end{pmatrix}$$

$$P=\begin{pmatrix} A & 0 & 0 & 0 \\ AB & A & 0 & 0 \\ AB^2 & AB & A & 0 \\ AB^3 & AB^2 & AB & A \end{pmatrix}$$

$$\{\mathrm{x1}[n] \to \lambda^n \mathrm{x1}[0]\}$$

$$\{\mathrm{x2}[n] \to \lambda^{-1+n}(-A(1+B)n\mathrm{x1}[0] + \lambda\mathrm{x2}[0])\}$$

$$\{\mathrm{x3}[n]$$

$$\{\mathrm{x4}[n]$$

$$\to \frac{1}{6}\lambda^{-3+n}$$

$$\times (-A^3(1+B)^3n(2-3n+n^2)\mathrm{x1}[0] + 3A^2(1+B)^2(-1+n)n\lambda(2B\mathrm{x1}[0] + \mathrm{x2}[0])$$

Monotonic Convergence (MC) of the Vector Norm#

- Worried about "bad learning behaviour", many authors have utilised a more robust form of convergence: leading to more stringent conditions on {K,τ;ν}.
  - Suppose we insist $(\mathbf{y}_{j+1})^{T*}(\mathbf{y}_{j+1}) = \lambda\times(\mathbf{y}_j)^{T*}(\mathbf{y})$ where T* means conjugate transpose.
  - It follows that $(\mathbf{y}_j)^{T*}[\lambda\mathbf{I}-\mathbf{A}^{T*}\mathbf{A}]\mathbf{y}_j=\mathbf{0}$, which is an eigenvalue equation.
  - "Monotonic convergence" then amounts to finding conditions on internal gains such that all eigenvalues of $\mathbf{A}^{T*}\mathbf{A}$ have modulus@ less than 1.

For brevity, let $\mathbf{y}_{j+1} = \mathbf{A}\,\mathbf{y}_j$.

- Let $\mathbf{S} \equiv \mathbf{A}^T\mathbf{A}$. The inequality $|\mathbf{Au}| \leq \sigma|\mathbf{u}|$ where σ is maximum |λ| of Eigenvalues[S] implies monotonic convergence of the vector norm* $|\mathbf{u}|=\mathrm{Sqrt}[\mathbf{u}^T.\mathbf{u}]$
- The ILC system behaves according to the eigenvalues of **A** with {K,τ;ν} chosen from the condition |Eigenvalues[S(K,τ;ν)]|<1
- For this subset {K,τ;ν}, convergence is stronger; and so the learning transients are reduced or eliminated.

2-norm or Euclidean norm or Euclidean length
@ The condition yields real eigenvalues, but does not prohibit them from being negative if the learning gain is pushed too hard.

Graphical representation of the influence of λ[S] on λ[A]

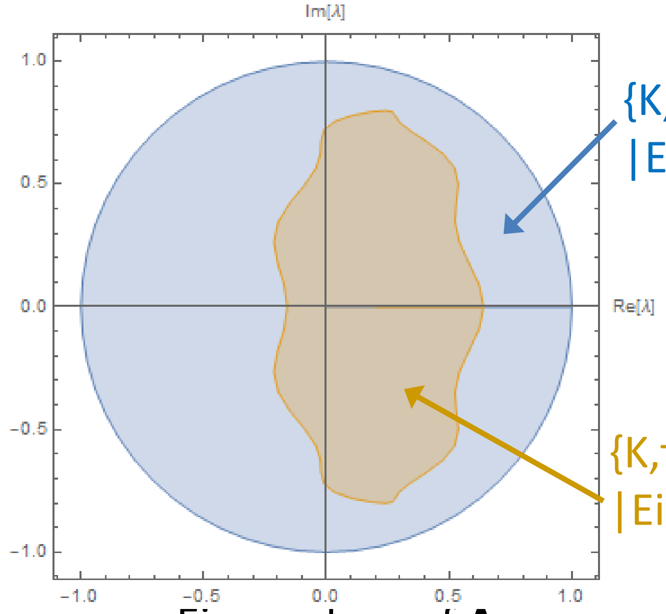

Eigenvalues of **A**

{K,τ,ν} chosen from
|Eigenvalues[A(K,τ,ν)]|<1

Noncausal Learning

{K,τ,ν} chosen from
|Eigenvalues[S(K,τ,ν)]|<1

Power series coefficients

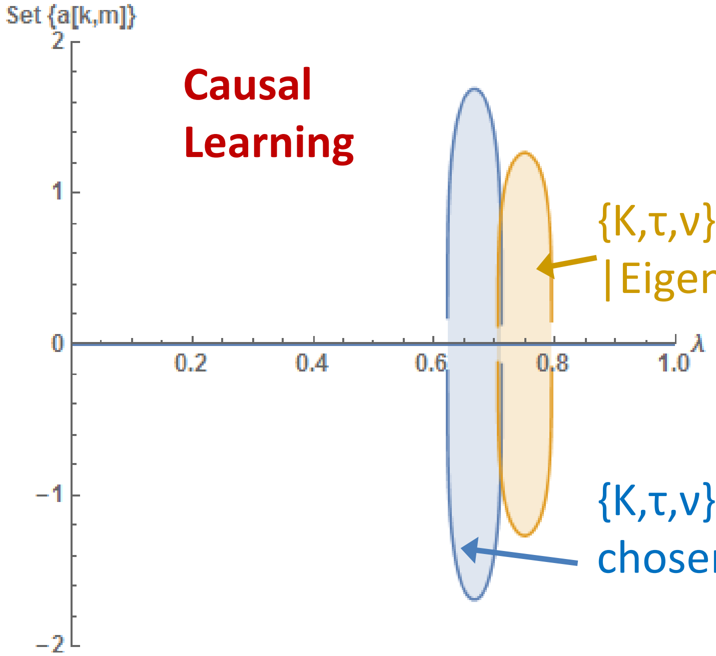

Causal Learning

{K,τ,ν} chosen from
|Eigenvalues[S(K,τ,ν)]|<1

{K,τ,ν} underdetermined
chosen from single Eigenvalue[A(K,τ,ν)]|<1

In both cases, subject to the constraint of system stability without ILC {K,τ;0}

- **Q: What other solutions?**
  - **in the subspace between the AC and MC conditions**
  - **Inside the space of MC conditions**
  - **away from fixed point**
  - **that couple within-trial behavior to iteration index**
- **A: Experiment by direct iteration, see what happens…**

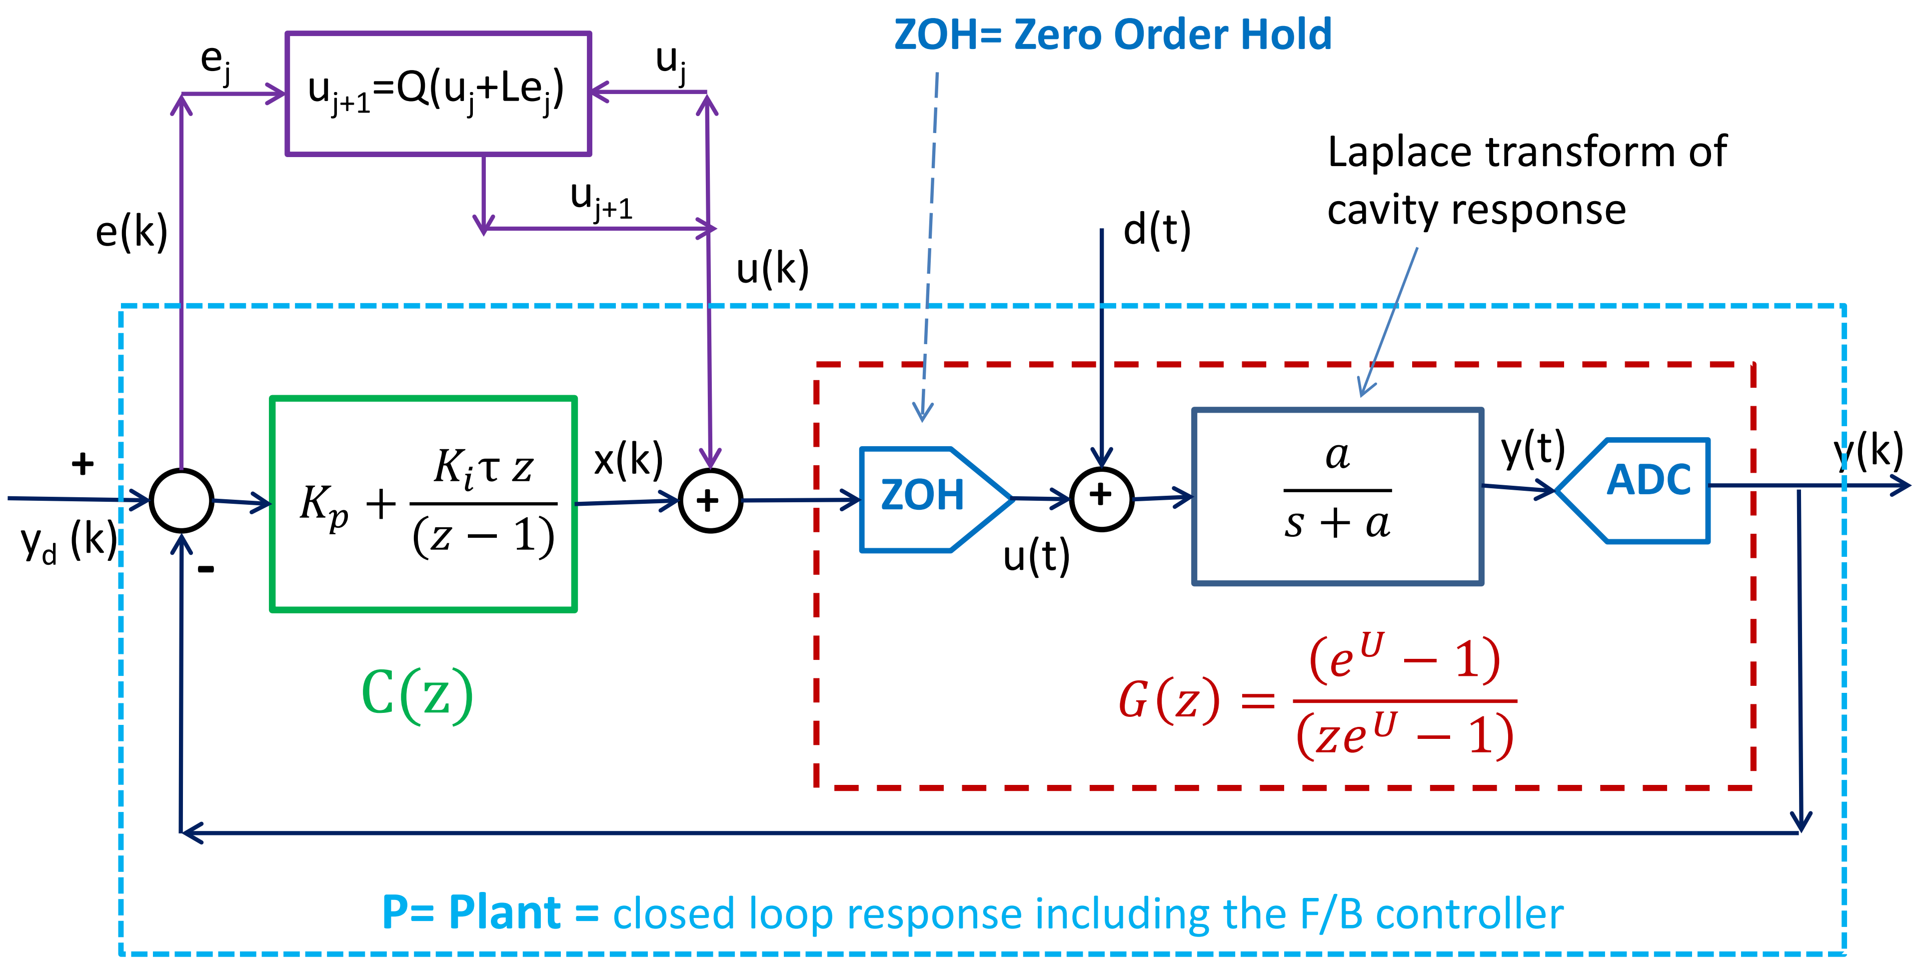

$$P(z)=\frac{G(z)}{1+C(z)G(z)}$$

$$y(z)=P(z)u(z)$$

Laplace & Matrix representation of the Plant

Z-domain Closed Loop Gain Function
Proportional Control K_p

Sample rate $\rho_s = 1/\tau_s$, where τ_s is sample period.

Cavity time constant is $\tau_c = 1/a$

$U = a\tau_s = 1/(\tau_c\, \rho_s) > 0$

$$P(z) = \frac{-1+e^{U}}{-1-Kp+e^{U}Kp+e^{U}z}$$

P(z) can be written: $\dfrac{A}{z-B}$

$$A = 1 - e^{-U} \geq 0 \qquad B = e^{-U}(1+Kp) - Kp$$

Low dimension (10x10) matrix representation of the Plant Operator

$$\mathbf{P} = \begin{pmatrix}
A & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
AB & A & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
AB^2 & AB & A & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
AB^3 & AB^2 & AB & A & 0 & 0 & 0 & 0 & 0 & 0 \\
AB^4 & AB^3 & AB^2 & AB & A & 0 & 0 & 0 & 0 & 0 \\
AB^5 & AB^4 & AB^3 & AB^2 & AB & A & 0 & 0 & 0 & 0 \\
AB^6 & AB^5 & AB^4 & AB^3 & AB^2 & AB & A & 0 & 0 & 0 \\
AB^7 & AB^6 & AB^5 & AB^4 & AB^3 & AB^2 & AB & A & 0 & 0 \\
AB^8 & AB^7 & AB^6 & AB^5 & AB^4 & AB^3 & AB^2 & AB & A & 0 \\
AB^9 & AB^8 & AB^7 & AB^6 & AB^5 & AB^4 & AB^3 & AB^2 & AB & A
\end{pmatrix}$$

I test for convergence w.r.t. matrix dimension; and typically use (100x100)

Matrix representation of example Learning functions

(banded diagonal forms)

3-term look-ahead $\mathbf{L}=\nu(I+\uparrow+\uparrow)$

$$\nu\begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

3-term look-back $\mathbf{L}=\nu(I+\downarrow+\downarrow)$

$$\nu\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \end{pmatrix}$$

The matrix dimension (N) is the trial length = (trial duration)/(sampling period) = T/τ_s
In numerical experiments we have to form high powers (M) of the matrix A=Q(I-PL) of dimension N^2, or order $(N^2)^M$ operations. M runs into hundreds.

As a function of matrix size, eigenvalues converges when dimension is order of magnitude larger than number of non-zero terms in a column (or row) of the banded diagonal.
Typically I use matrix dimension (100x100) to represent the 3-term learning functions

Example: domains of monotonic convergence

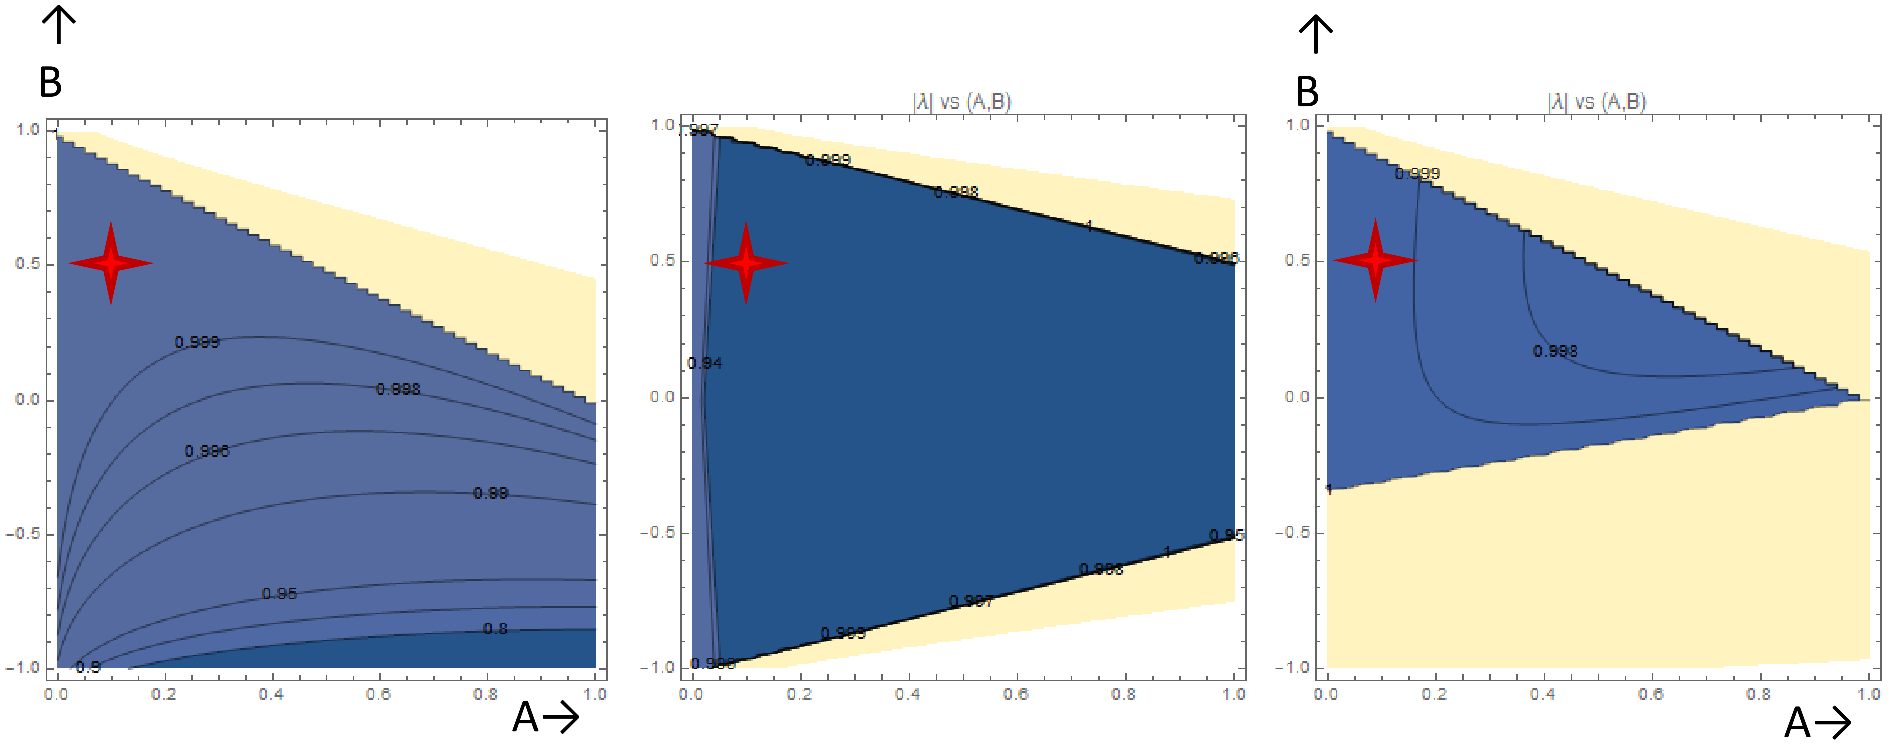

2-term look-back
L=(I+↓)

1-term look-back
L=I

2-term look-ahead
L=(I+↑)

Blue=monotonic convergent
Yellow=not monotonic

L = Identity matrix

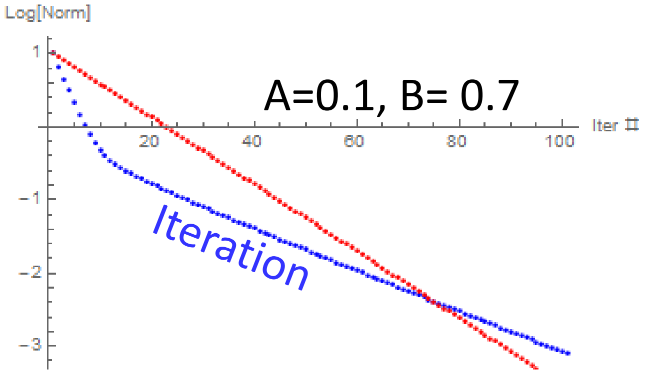

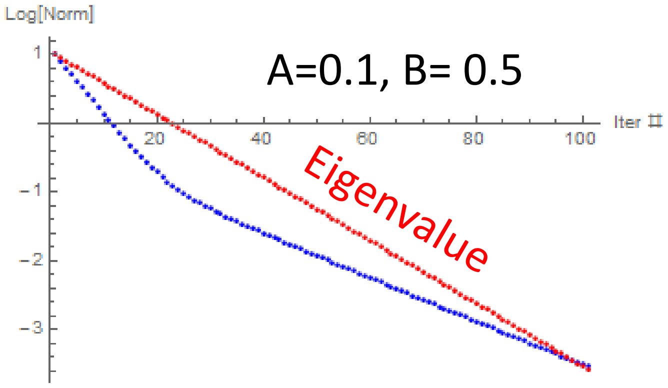

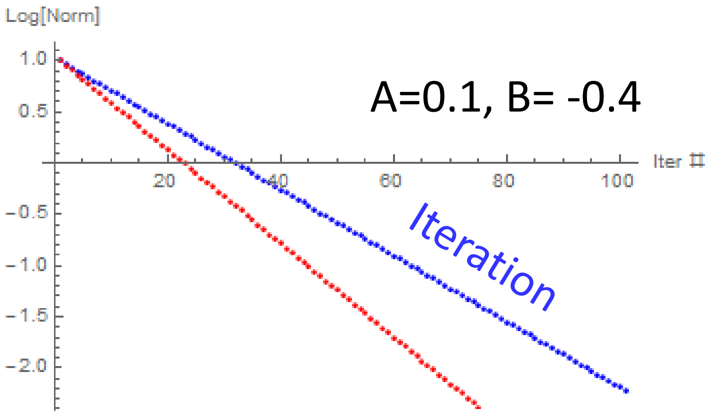

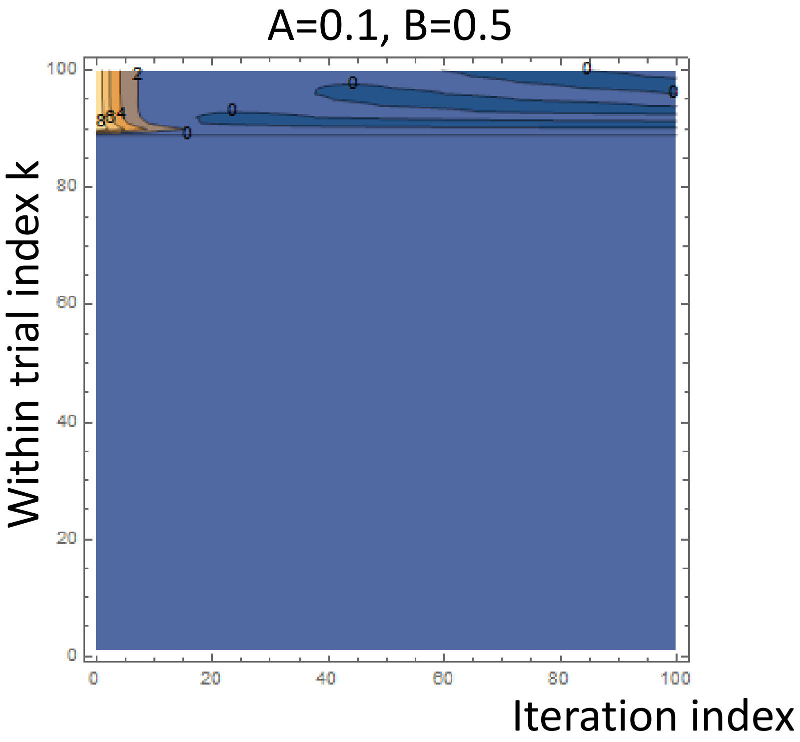

Behaves as expected:

Approx linear decrement of log[norm(**u**)]

No bad learning transients for these conditions

2-term look ahead; L=(I+↑)

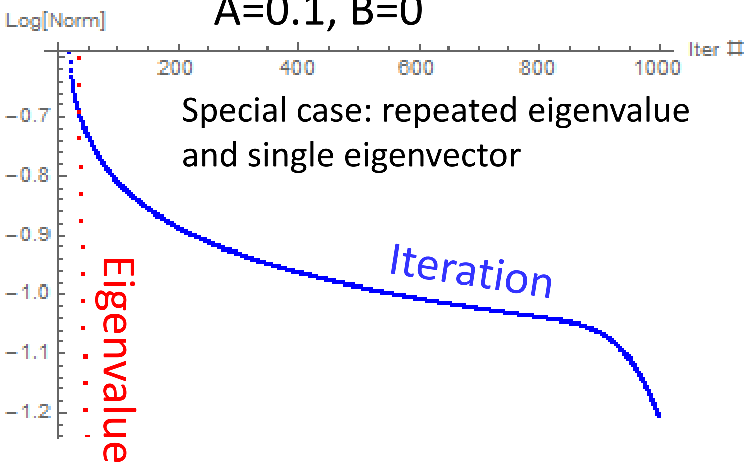

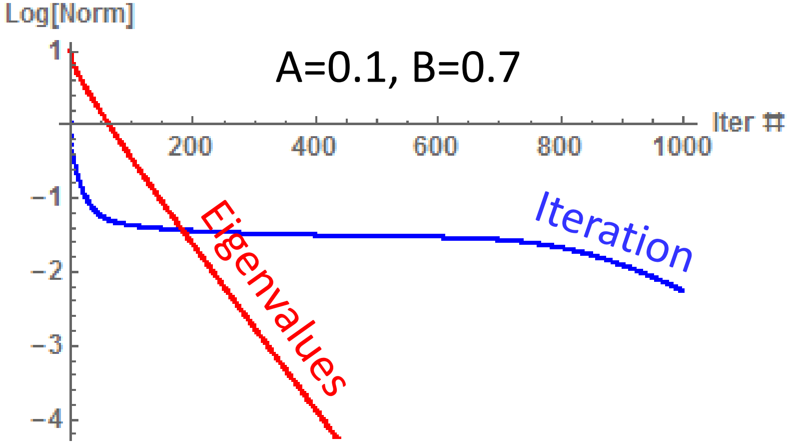

Wave propagation

δ-function disturbance in last element of initial vector

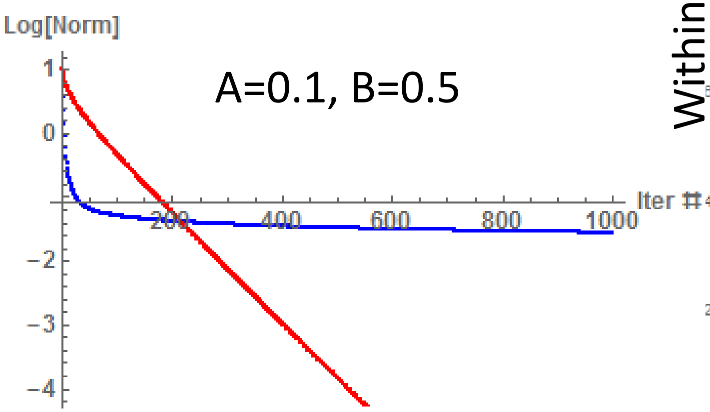

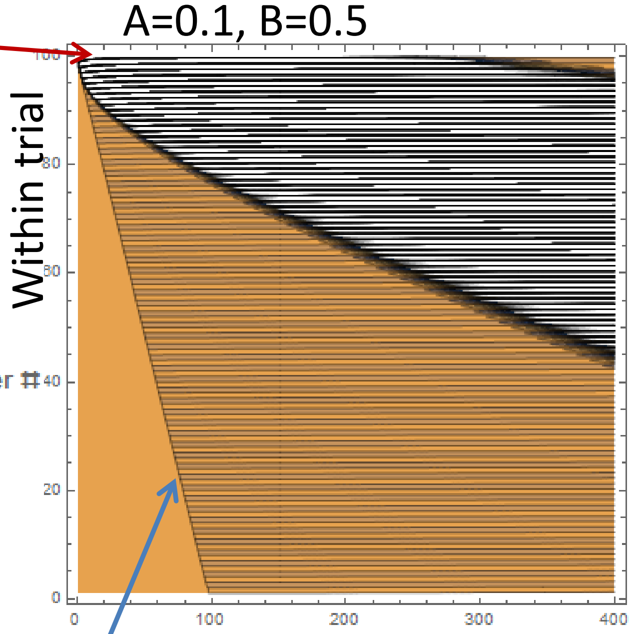

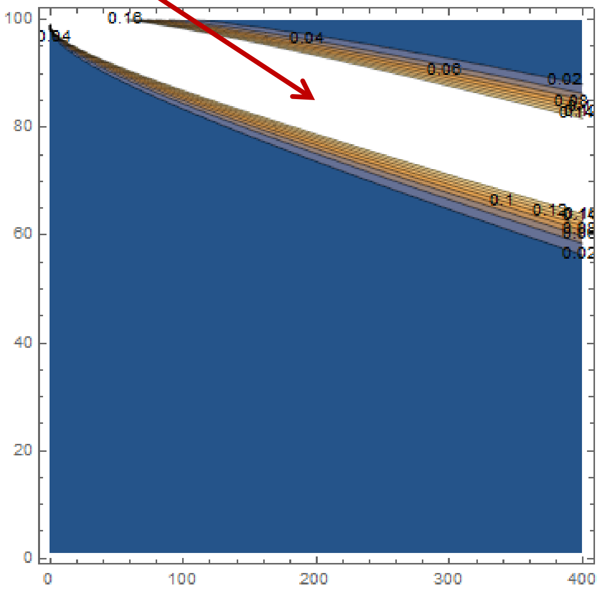

Iterates

Algorithm propagation

Wave travels from “late” toward “early” k

2-term look ahead; L=(I+↑)

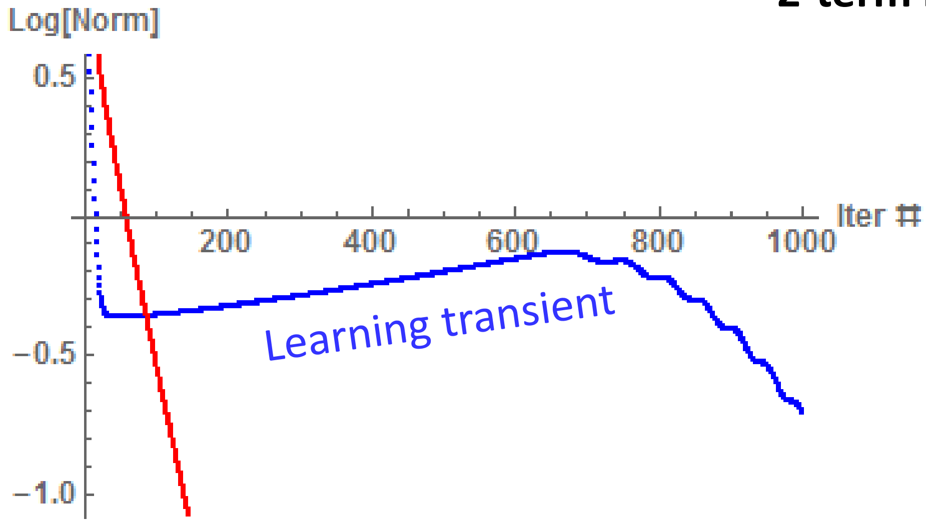

A=0.1, B= -0.35

ILC Algorithm propagation

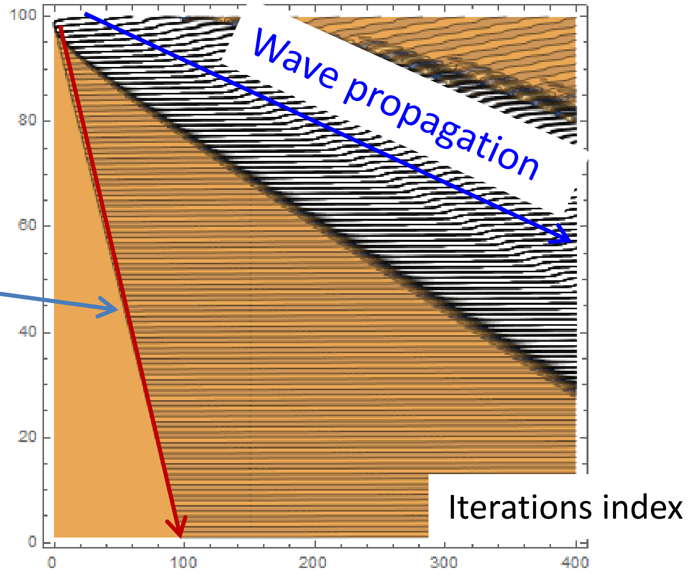

A=0.1, B= -0.5

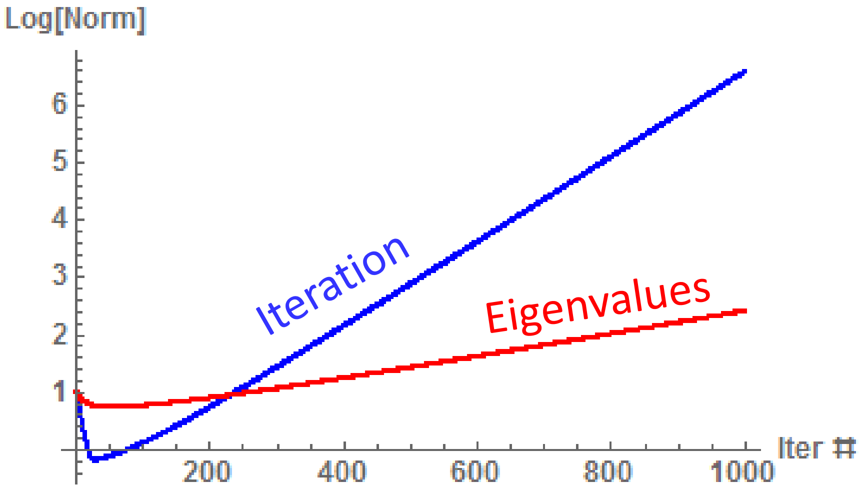

NOT a learning transient, because|λ|>1

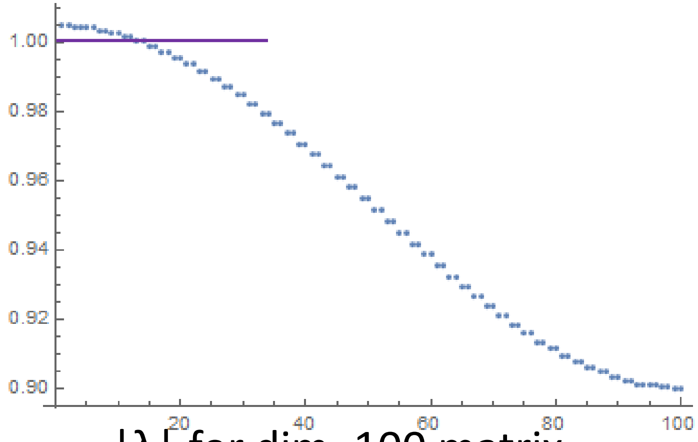

|λ| for dim=100 matrix

2-term look back L=(I+↓)

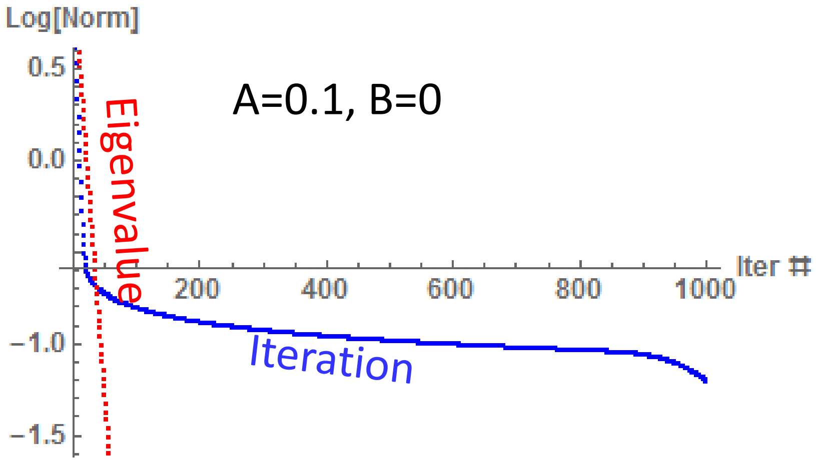

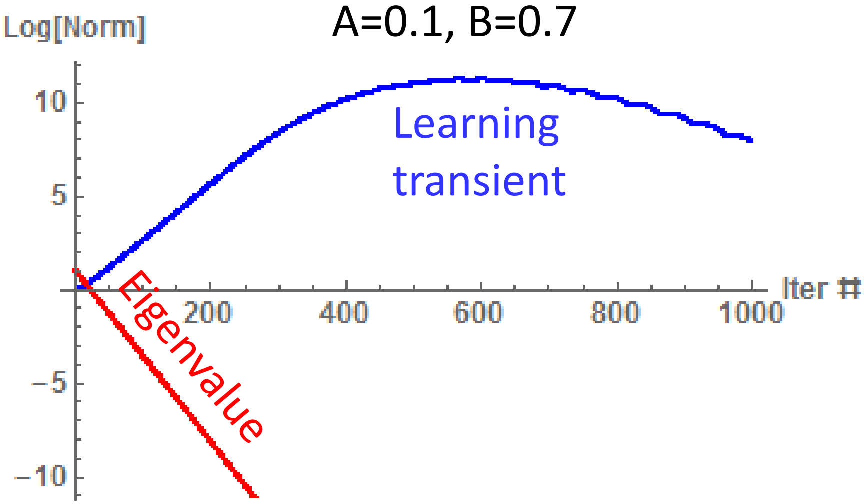

A=0.1, B=0.5

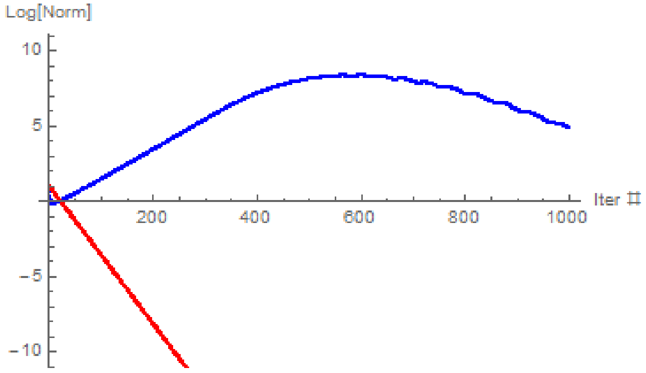

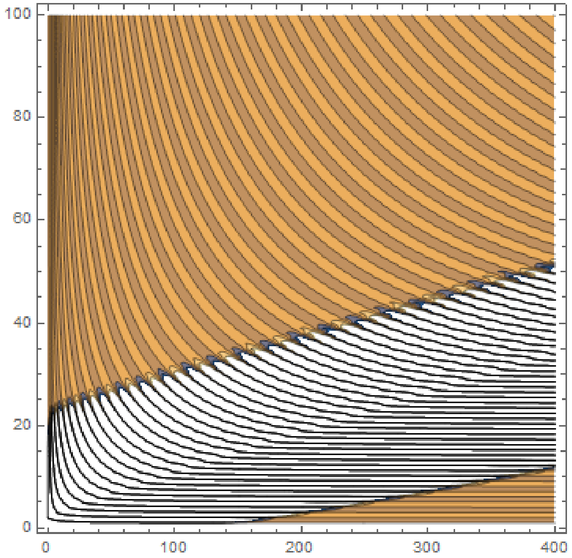

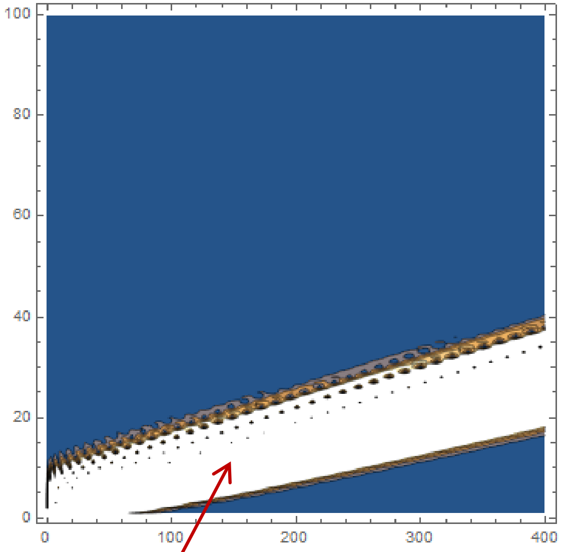

Within trial

Iterates

Wave propagation

Initial δ-function disturbance
in the first element

Wave travels from “early” toward “late”

2-term look back L=I+↓; A=0.1, B=0.5

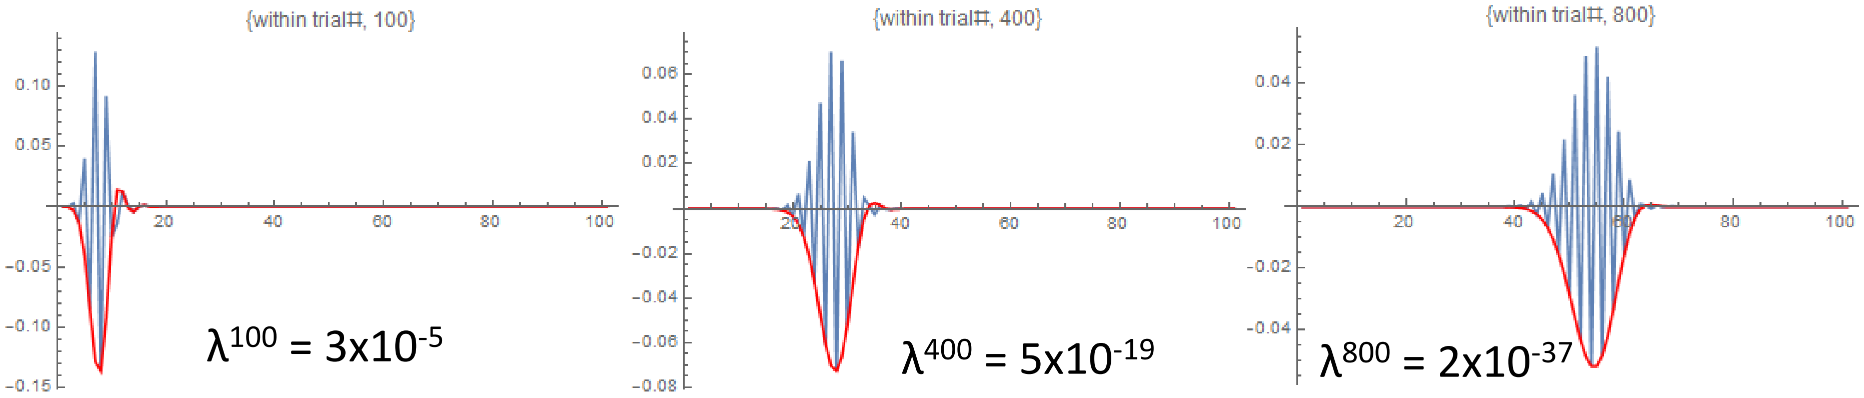

2-term look ahead L=I+↑; A=0.1, B=0.5

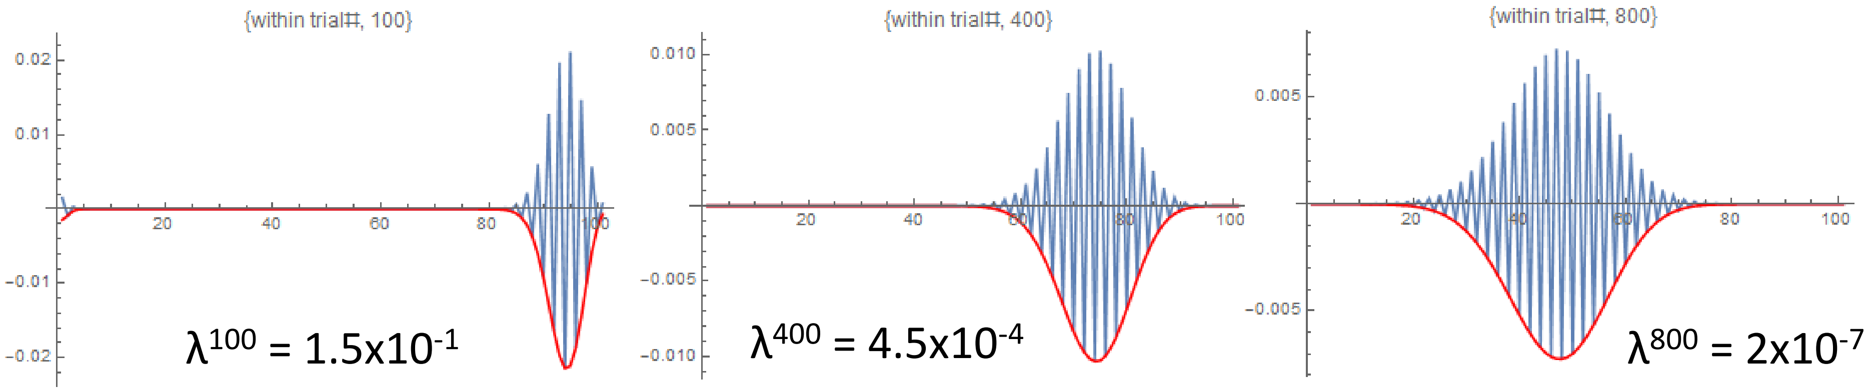

- We are looking for solutions that survive hundreds or more iterations.
- Definition of a wave:
  - An object that moves with little or no change of shape at an identifiable speed.
- Yes “waves” look like a candidate.

Q: Does $\mathbf{u}_{n+1} = \mathbf{A}^n \mathbf{u}_1$ have wave-like solutions for large n?

A: Experimentally, a clear "Yes".

Mathematically, must satisfy a wave equation – something like

$$\left[\frac{\partial \boldsymbol{u}}{\partial n} + c\,\frac{\partial \boldsymbol{u}}{\partial k}\right] = 0$$

- .*n* = iteration index, k = row index of the column vector **u**, and c = wave velocity.
- Have to find analogs of the differential operators
- Have to find correct waveform
  - Either from experimental data or a guess
  - [Of course, it is possible there are no such waveforms.]

Discrete Analogs of Differential Operators

Waveform $u[k,n] == \mathrm{Cos}[\frac{2\pi k}{2}]\, S[\frac{k-c\,n}{\sigma}]$

S="shape"

Cos[πk] carrier is known from 2-term experiments

Continuous	Discrete
$\frac{\partial \boldsymbol{u}}{\partial n}$	$\frac{\boldsymbol{\Delta u}}{\boldsymbol{\Delta n}} == (\boldsymbol{A}-\mathbf{I})\boldsymbol{u} == (\boldsymbol{Q}(\mathbf{I}-\mathbf{PL})-\mathbf{I})\boldsymbol{u} \;\; \rightarrow -\boldsymbol{PLu}$ if **Q=I**
$\frac{\partial \boldsymbol{u}}{\partial k}$	In element form Δu/Δk = u[k+1,n]-u[k,n] = Cos[π(k+1)]S[k+1,n]-Cos[πk]S[k,n]= (-↑-I)u[k,n]

for this particular waveform

$$\boxed{\frac{\boldsymbol{\Delta u}}{\boldsymbol{\Delta k}} \rightarrow -(\uparrow + \boldsymbol{I})\boldsymbol{u}}$$

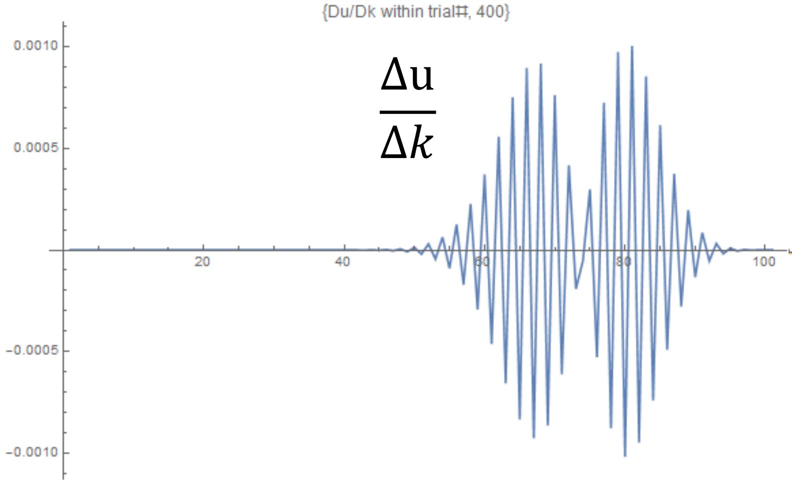

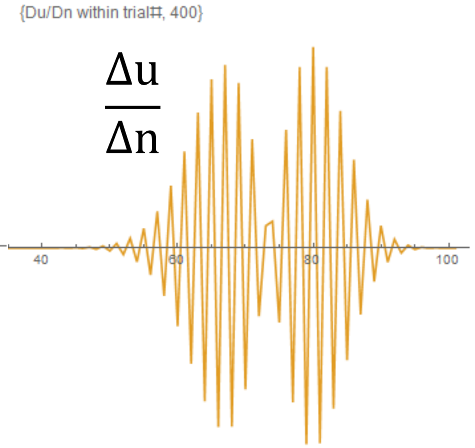

2-Term look ahead learning function; A=0.1, B=0.5

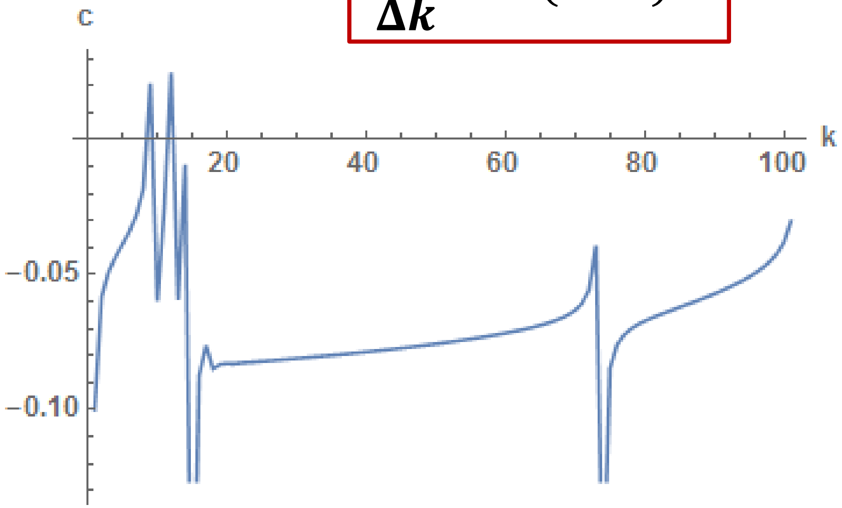

Element by element $c == \frac{(A-\mathrm{I})u}{(\uparrow + \mathrm{I})u}$

For 2-term learning,
we must find the self-consistent waveform **u** and wave velocity *c* that satisfies

$$(\boldsymbol{A}-\mathbf{I})\boldsymbol{u} == c(\uparrow+\mathbf{I})\boldsymbol{u}$$

We know $u[k,n] == (-1)^k S[X]$

LHS→ function; RHS→ derivative

Therefore shape S[X]=Exp[f(X)]

The wave neither grows nor decays; therefore f(x)→ - ½ X^2

Insert $u[k,n] == (-1)^k S[-X^2/2]/(2\sigma)$ into the wave equation,
set cn=k_0 the wave centre $X=\frac{k-k0}{\sigma}$

- Compare each row, *k,* → *N* equations for speed *c* with 1 free parameter, σ.
- If all equations converge on same *c* value as σ is varied, then we have a self-consistent solution. (N is the matrix size)

Wave Solutions to Iterative Learning

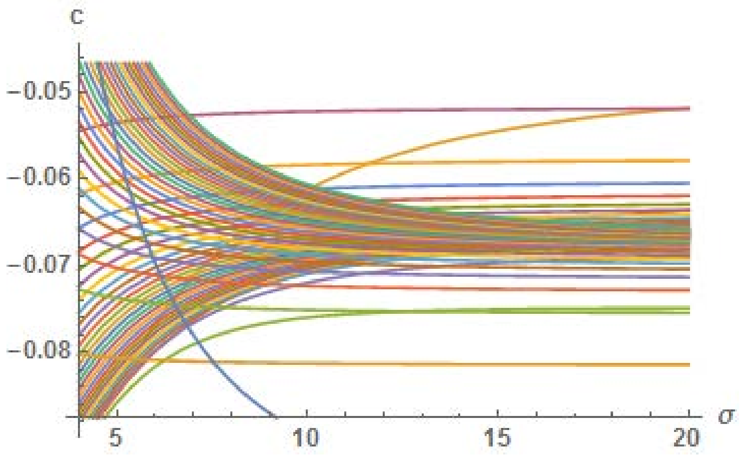

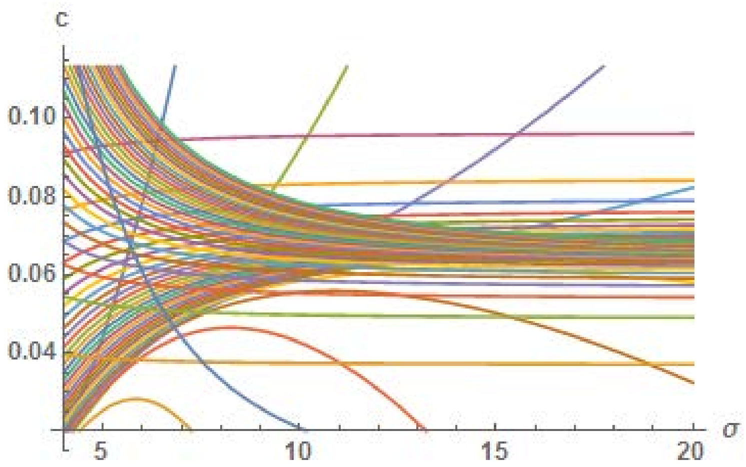

Converges to common speed *c*

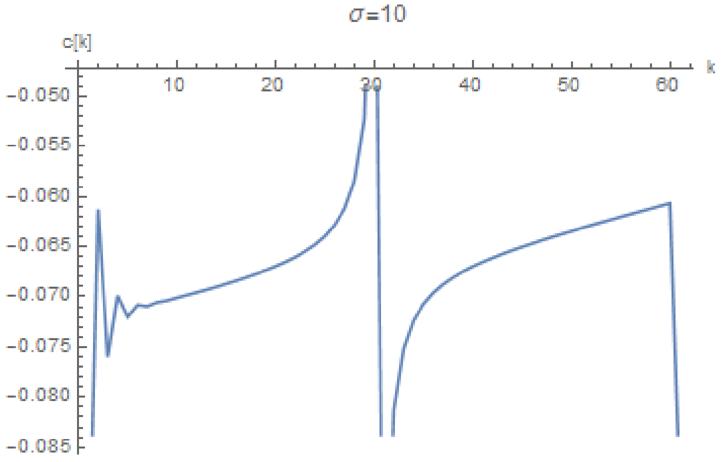

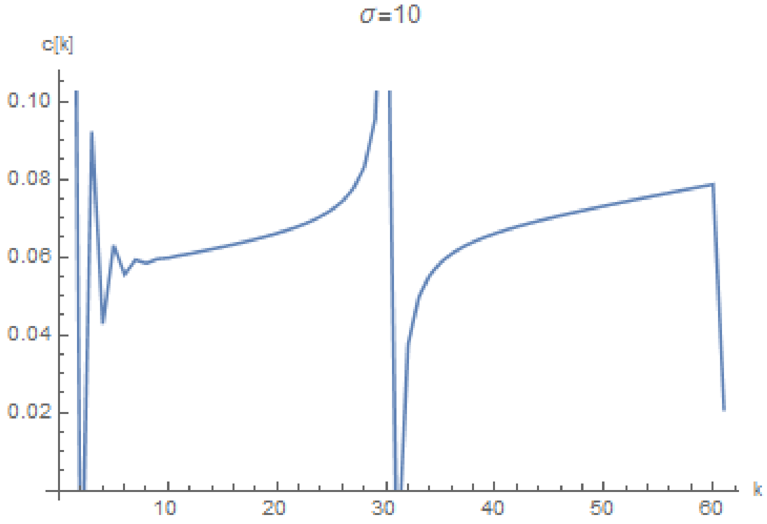

2-term Look-Ahead (I+↑)
A=0.1,B=0.5

2-term Look-Back (I+↓)
A=0.1,B=0.5

(A,B) Values are inside domain of monotonic convergence

Discrete Analogs of ∂u/∂k differential Operator

- The learning function sets the period and form of the carrier
- For Identity and 2-term learning, the wave is u[k,n]=Cos[2π(k/2)]S[k,n]
- For 3-term learning, wave is u[k,n]={Cos[2π(k/3)]+q1 Cos[2π(k±1)/3]} S[k,n]
- For 4-term learning, wave is {Cos[2π(k/4)]+q1 Cos[2π(k±1)/4]+q2 Cos[2π(k±2)/4]}S[k,n]
- And so on…
- Take minus sign for look-back; plus for look-ahead.
- In element form $\Delta u/\Delta k|_k$ =u[k+1,n]-u[k,n]

3-term look-back $\quad \Delta u/\Delta k|_k = -\frac{1}{2}S[k,n]\left((2+\uparrow-\mathrm{q1}-2\uparrow\mathrm{q1})\,\mathrm{Cos}\left[\frac{2k\pi}{3}\right]+\sqrt{3}(\uparrow+\mathrm{q1})\,\mathrm{Sin}\left[\frac{2k\pi}{3}\right]\right)$

In operator form $\quad \boxed{\frac{\Delta u}{\Delta k}\to[r\uparrow-\mathbf{I}]u}$ $\quad r_{k,k} = \frac{(1-2\mathrm{q1})\,\mathrm{Cos}\left[\frac{2k\pi}{3}\right]+\sqrt{3}\,\mathrm{Sin}\left[\frac{2k\pi}{3}\right]}{(-2+\mathrm{q1})\,\mathrm{Cos}\left[\frac{2k\pi}{3}\right]-\sqrt{3}\mathrm{q1}\,\mathrm{Sin}\left[\frac{2k\pi}{3}\right]}$ $\quad r_{k,j}=0$, $k\neq j$

4-term look-back

$\Delta u/\Delta k|_k = S[k,n]\left((-1+\uparrow\mathrm{q1}+\mathrm{q2})\,\mathrm{Cos}\left[\frac{k\pi}{2}\right]-(\uparrow+\mathrm{q1}-\uparrow\mathrm{q2})\,\mathrm{Sin}\left[\frac{k\pi}{2}\right]\right)$

In operator form $\quad \boxed{\frac{\Delta u}{\Delta k}\to[r\uparrow-\mathbf{I}]u}$ $\quad r_{k,k} = -\frac{\mathrm{q1}\,\mathrm{Cos}\left[\frac{k\pi}{2}\right]+(-1+\mathrm{q2})\,\mathrm{Sin}\left[\frac{k\pi}{2}\right]}{(-1+\mathrm{q2})\,\mathrm{Cos}\left[\frac{k\pi}{2}\right]-\mathrm{q1}\,\mathrm{Sin}\left[\frac{k\pi}{2}\right]}$ $\quad r_{k,j}=0$, $k\neq j$

Conclusions: ILC Gone Wild

- ***Well Known***
- If convergent, ILC maps iterate to their fixed points (FP)
- If **Q=I**, then FP=0. If **Q≠I**, then FP≠0 and residual error **e≠0**
- Near the FP, noncausal learning has eigensolutions such that $\mathbf{A}^n\mathbf{e}_k=\lambda^n{}_k\mathbf{e}_k$
- Causal and noncausal learning both display transients for {K,τ;ν} outside the MC-domain
- ***Less Well Known***
- Causal learning has eigensolutions $\mathbf{A}^n\mathbf{e}_k=\lambda^n\mathbf{e}_k(n)$ with $\mathbf{e}_k(n)=\lambda^{-k}\,\Sigma_m\ a[m,k]n^m$
  - eigensolution #k starts to converge for n>k
- Gains {K,τ;ν} chosen from MC condition $|\lambda(A^TA)|<1$ influence λ_k and λ & a[k,m]
  - Subject to **P** alone is stable for {K,τ;0}

In the companion paper "**ILC – Deep Dive**", or "**Testing the Tests**" author
Applied: Z-dom MC test & λ-dom MC test & λ-dom AC test to causal and noncaual learning functions
Found: both MC tests & single AC-test agree with results of direct iteration of ILC equations.
Found: Z-dom MC test & λ-dom MC conditions generate the same range of stable parameters {K,τ;ν}.

But direct iteration reveals the λ-domain AC-test to be almost useless – because, for parameters {K,τ;ν} outside the MC domain, transients may reach astronomically large values before finally converging

Conclusions: ILC Gone Wild

- WAVES FOUND IN 2016
- When **Q=I** and **L** is the diagonal-band style of learning function $\mathbf{L}=\nu_0\mathbf{I}+\Sigma_q\nu_q\boldsymbol{\uparrow}^q+\Sigma_p\nu_p\boldsymbol{\downarrow}^p$
- There may also be soliton wave-type solutions S(*k*-*c.n*) existing <u>inside the MC domain</u>. The waves* have carrier close to the Nyquist frequency $\rho_s/2$
- The waves obey the equation **(A-I)u** = c**(Δ/Δk)u** where the wave shape S and speed c must be found self-consistently.
- Whereas the phase velocity of propagation of the ILC algorithms is c=k/n=1,
- The group velocity of the wave c= **Δk/Δu** <<1
- The waves may be suppressed by a low pass filter **Q**, but at the cost of introducing residual errors, and changing the eigensolutions.

* If the learning function contains $\uparrow^q$ or $\downarrow^q$ the carrier frequency is ρ_s/q

Z-domain Stability and Convergence Within a Trial

- Within a single trial, response of plant P evolves according to stability of P without Iterative Leaning Control.
- Sampled system → use Z-transform (Laplace transform for discrete time)
  - Either: (i) examine directly poles and zeroes; or (ii) plot poles of the closed loop transfer function in complex plane.

- Stability within the trial including the internal dynamics introduced by the learned time function u is addressed by the Z-transform response of the plant with ILC
- Because it is about stability within trials, it makes no direct connection to iteration eigenvalues.
- However, it can tell us if system settles down, and if so then the iterations must have converged.
  - And yields conditions on {K,τ;ν} for Monotonic Convergence

In the extended Z-domain of trials and iterations, the iteration equation becomes

$u_{j+1}(z)=T(z)u_j(z)$ with $T(z)= Q(z)[1 - z^m P(z)L(z)]$ and* $m \equiv 1$.
Let $z=e^{i\theta}$. The inequality Supremum$[T(e^{i\theta})]<1$ implies conditions on the gain set {K,τ;ν} very similar (often identical) to the Monotonic Converge condition above.

In practise: form the locus of T with $z=e^{i\theta}$; look for largest values → {K,τ;ν} constraints.

(NOTE: there is no Z-domain analog of the Asymptotic Convergence condition.)

*z^m accounts for the “lift” that must be applied to the matrix representation of P to compensate for the delay introduced by the zero-order-hold.

- R.W. Longman: IEEE Trans. Circuits & Systems, 2002, Vol.49, No.6, 753-767
- “consecutive waves of convergence along the column vectors” explanation of “bad” causal learning behaviour
- Is almost (but NOT quite) the explanation

ILC Algorithm propagation

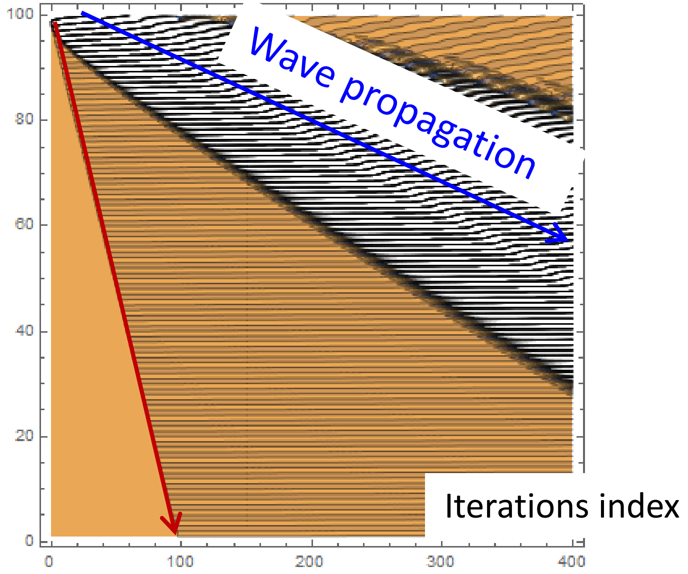

The phase velocity of the disturbance corresponds to the “waves of convergence along the column vectors”
But the “real” disturbance moves as a wave at the group velocity.