




Data-driven transient growth analysis

Zhicheng Kai ^{*}, Peter Frame , and Aaron Towne *Department of Mechanical Engineering, University of Michigan, Ann Arbor, Michigan, USA*

(Received 5 June 2025; accepted 13 January 2026; published 20 March 2026)

The transient growth of disturbances made possible by the non-normality of the linearized Navier-Stokes equations plays an important role in bypass transition for many shear flows. Transient growth is typically quantified by the maximum energy growth among all possible initial disturbances, which is given by the largest squared singular value of the matrix exponential of the linearized Navier-Stokes operator. In this paper, we propose a data-driven approach to studying transient growth wherein we calculate optimal initial conditions, the resulting responses, and the corresponding energy growth directly from flow data. Mathematically, this is accomplished by optimizing the growth over linear combinations of input and output data pairs. We also introduce a regularization to mitigate the sensitivity to noisy measurements and unwanted nonlinearity. The data-driven method simplifies and broadens the application of transient growth analysis: it removes the burden of writing a new code or linearizing an existing one, alleviates the computational expense for large problems, eliminates the challenge of obtaining a well-posed spatial propagator for spatial growth analyses, and enables the direct application of transient growth analysis to experimental data. We validate the data-driven method using a linearized Ginzburg-Landau model problem and Poiseuille flow corrupted by process and measurement noise and obtain good agreement between the data-driven and the standard operator-based results. We then apply the method to study the spatial transient growth of disturbances in a transitional boundary layer using data from the Johns Hopkins Turbulence Database. Our method successfully identifies the optimal output response and provides plausible estimates of the transient spatial energy growth at various spanwise wave numbers.

DOI: [10.1103/mxyg-j917](https://doi.org/10.1103/mxyg-j917)

I. INTRODUCTION

The classical approach to hydrodynamic stability investigates the long-time asymptotic behavior of infinitesimal disturbances to the base flow of interest. This long-time behavior is determined by the sign of the eigenvalues of the linearized Navier-Stokes (LNS) operator: if there is an eigenvalue with a positive real part, infinitesimal disturbances grow exponentially in the long term; otherwise, they decay. Of course, physical disturbances to the flow are not infinitesimal, and the assumption underlying the relevance of infinitesimal disturbances is that realistic disturbances are initially small enough that nonlinear effects are negligible. This approach to determining stability, usually referred to as modal stability theory, has many successes. For instance, it gives an accurate analytic prediction for the critical Rayleigh number in Rayleigh-Bénard convection [1,2] and for the critical Taylor number in Taylor-Couette flow [3]. However, modal stability theory predicts stability when experiments exhibit turbulence for many wall-bounded shear flows, including many of the most canonical problems in stability theory, such as pipe and Couette flows [4–6]. The phenomenon

^{*}Contact author: kaizz@umich.edu

of transient growth [5,7–9] offers a resolution to this issue. The non-normality of the evolution matrix can cause the disturbances to grow temporally, even in an asymptotically decaying system. This transient growth is usually quantified by the optimal growth G^{opt} and the associated modes representing the optimal initial condition and response. G^{opt} is the maximum growth in the kinetic energy caused by any initial condition. The theory of transient growth has provided physical insight into previously unexplained instabilities in many shear flows in both the temporal [5] and spatial settings [10].

The standard method for calculating G^{opt} and the associated modes requires access to the LNS operator, but this operator is often difficult to obtain. Deriving the linearized equations and writing a stability code for a new problem can be tedious, as is extracting the linear operator from an existing nonlinear computational fluid dynamics code [11,12]. If spatial (rather than temporal) growth is of interest, great care must be taken to obtain a stable spatial-evolution operator [13,14]. Once the necessary operator is in hand, computing the optimal growth can be computationally costly, especially for problems that lack homogeneous directions that can be used to decouple individual Fourier modes and reduce the size of the associated operators.

In this paper, we aim to simplify and broaden the application of transient growth analysis by developing an algorithm that approximates G^{opt} using data. The method uses a set of initial- and final-state pairs and assumes that there is an (unknown) linear operator relating the two. With this assumption, the final state resulting from any linear combination of the initial states in the data is the same linear combination of the corresponding final states. The algorithm approximates G^{opt} as the maximum energy ratio of the initial and final states over all such linear combinations. The method also includes a regularization to mitigate the effect of noise.

The data-driven approach removes the burden of writing new code or linearizing an existing one. It also substantially alleviates the computational expense: the proposed method is roughly as expensive as obtaining proper orthogonal decomposition (POD) modes [15]. Finally, it enables the direct application of transient growth analysis to experimental data, which is not possible with the standard operator-based approach.

Our approach differs in a critical way from the work of Ref. [16], another recently proposed approach to estimating transient growth from data. That work [16] uses dynamic mode decomposition (DMD) to approximate the one-space-step evolution operator, finds the disturbance that grows the most after one spatial step, and then evolves this disturbance forward in the streamwise direction (x) using the same evolution operator at every x location. The optimal transient growth curve, however, is not the energy of any single disturbance as it evolves downstream. Instead, it is the envelope of these energies. The disturbance that grows most rapidly initially is not the disturbance that maximizes the growth at later times. Additionally, the operator that maps a disturbance Δx forward in space changes as a function of the x location due to the nonparallel nature of flows investigated with spatial stability. Our approach resolves both issues.

We first test our method on a linearized Ginzburg-Landau system. We compare the results of our data-driven transient growth algorithm to those obtained using the standard (operator-based) approach and find good agreement even in the presence of significant measurement and process noise. We also use the Ginzburg-Landau system to inform our choice for a parameter that controls the strength of the regularization.

Finally, we apply the method to transitional-boundary-layer data from the Johns Hopkins Turbulence Database (JHTDB) [17–19]. In this application, we intend to identify spatial energy growth—the growth of disturbances as they evolve in the streamwise direction. While performing a traditional transient growth analysis is challenging for this problem—the problem dimension is large, and the spatial LNS operator is not readily available—the data-driven method scales favorably to large problems and does not require the LNS operator, so it can be readily applied. The G^{opt} results given by the method are in a reasonable range relative to past studies [20,21] on a Blasius boundary layer but are sensitive to a parameter used in the regularization, and are thus uncertain. On the other hand, the output modes we obtain are similar to those of Refs. [20,21]. In addition, the output modes obtained by the data-driven method at $\beta = 0$ displayed a two-peak structure, which

is indicative of modal growth, whereas, at nonzero β , the modes displayed a one-peak structure indicative of nonmodal growth.

The remainder of this paper is organized as follows. We first review the standard (operator-based) transient growth theory in Sec. II before introducing the data-driven method in the context of temporal stability in Sec. III. In Sec. IV, we use the Ginzburg-Landau equation to verify the accuracy of the data-driven method, assess the impact of various parameters, and determine a recommended value of a key regularization parameter. In Sec. VI, we apply our method to study the spatial growth of disturbances in a transitional boundary layer, and we conclude the paper in Sec. VIII.

II. TRANSIENT GROWTH ANALYSIS

The linearized Navier-Stokes equations,

$$\frac{d}{dt}\mathbf{q} = \mathbf{A}\mathbf{q}, \quad (1)$$

govern the flow disturbance $\mathbf{q} = \mathbf{q}_{\text{total}} - \bar{\mathbf{q}} \in \mathbb{C}^n$, where $\mathbf{q}_{\text{total}}$ is the total flow state, $\bar{\mathbf{q}}$ is the base flow, and $\mathbf{A} \in \mathbb{C}^{n \times n}$ is the LNS operator. Here, the state and equations have already been spatially discretized. The solution of Eq. (1) is

$$\mathbf{q}_t = e^{\mathbf{A}t} \mathbf{q}_0 = \mathbf{M}_t \mathbf{q}_0, \quad (2)$$

where $\mathbf{q}_t \in \mathbb{C}^n$ is the disturbance \mathbf{q} at time t , and \mathbf{q}_0 is the initial disturbance. $\mathbf{M}_t = e^{\mathbf{A}t}$ is the time-evolution operator that advances the initial condition to time t . Equations (1) and (2) lay the foundation for discussing flow stability and energy growth within this paper for the temporal case. In later sections, we also apply the method to the spatial stability case. The eigenvalues of \mathbf{A} determine the asymptotic growth rate of the solution of linearized system (1). If all of the eigenvalues have negative real parts, the norm of the state always decays asymptotically, i.e., $\lim_{t \rightarrow \infty} \|\mathbf{q}_t\| = 0$. However, if any eigenvalue has a positive real part, the corresponding eigenmode will grow exponentially, and so will the disturbance. Thus, this approach, referred to as modal stability theory, reduces the question of stability to an eigenvalue problem. Modal stability, however, fails to predict the transition to turbulence observed in many shear flows, and the potential for large-scale transient growth offers an explanation.

The key insight is that, though a disturbance to a stable system will decay in the long term, this decay is not required to be monotonic. Instead, as long as the eigenvectors are nonorthogonal, the sum may grow temporarily, and this growth is referred to as transient growth. Indeed, the eigenvectors of the LNS operator in shear flows are often highly nonorthogonal.

Optimal growth, defined as the maximum energy growth at a given time t across all initial conditions [5], is often used to quantify transient growth. Mathematically, the optimal growth is expressed as

$$G^{\text{opt}}(t) = \max_{\mathbf{q}_0} \frac{\|\mathbf{q}_t\|^2}{\|\mathbf{q}_0\|^2} = \max_{\mathbf{q}_0} \frac{\|\mathbf{M}_t \mathbf{q}_0\|^2}{\|\mathbf{q}_0\|^2}. \quad (3)$$

The norm, representing the kinetic energy, can be expressed as a weighted 2-norm,

$$\|\mathbf{q}\|^2 = \mathbf{q}^* \mathbf{W} \mathbf{q}, \quad (4)$$

where \mathbf{W} is the weight matrix and $(\cdot)^*$ denotes the Hermitian transpose. The weight matrix can be used to account for nonuniform grids and to impose a desired definition of energy, e.g., a kinetic energy norm for incompressible flow or the Chu energy norm for compressible flows [22]. Using the Cholesky decomposition $\mathbf{W} = \mathbf{L}^* \mathbf{L}$, the maximized energy growth and its corresponding modes can be obtained from the singular value decomposition (SVD) [23]

$$\mathbf{L} \mathbf{M}_t \mathbf{L}^{-1} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^*. \quad (5)$$

The optimal growth is given by the largest squared singular value, i.e., $G^{\text{opt}}(t) = \sigma_1^2$, and the corresponding input and output modes are the first columns of $\mathbf{L}^{-1}\mathbf{U}$ and $\mathbf{L}^{-1}\mathbf{V}$, respectively.

III. DATA-DRIVEN TRANSIENT GROWTH ANALYSIS

In this section, we describe a data-driven approach to approximating the transient growth analysis described in the previous section. We formulate the baseline method in Sec. III A before introducing a regularization used to reduce the sensitivity to noisy data and presenting the overall algorithm and its computational cost scaling in Sec. III B.

A. Formulation

Given data comprised of a set of initial disturbances $\{\mathbf{q}_0^1, \mathbf{q}_0^2, \dots, \mathbf{q}_0^m\}$ and the corresponding evolved disturbances $\{\mathbf{q}_t^1, \mathbf{q}_t^2, \dots, \mathbf{q}_t^m\}$, where

$$\mathbf{q}_t^k = \mathbf{M}_t \mathbf{q}_0^k, \quad k \in [1, m], \quad (6)$$

we form the data matrices

$$\mathbf{Q}_0 = [\mathbf{q}_0^1, \mathbf{q}_0^2, \dots, \mathbf{q}_0^m] \in \mathbb{C}^{n \times m}, \quad \mathbf{Q}_t = [\mathbf{q}_t^1, \mathbf{q}_t^2, \dots, \mathbf{q}_t^m] \in \mathbb{C}^{n \times m}. \quad (7)$$

Recall that the state \mathbf{q} is defined as the perturbation to some base state $\bar{\mathbf{q}}$ such that it obeys the linear system in Eq. (1).

Next, consider a new initial condition constructed as a linear combination of the set of initial disturbances,

$$\mathbf{q}_0 = \mathbf{Q}_0 \psi, \quad (8)$$

where $\psi \in \mathbb{C}^m$ is a vector of expansion coefficients. Since the system is linear, the response at future times to this initial disturbance is obtained from substituting Eq. (8) into Eq. (6):

$$\mathbf{q}_t = \mathbf{M}_t \mathbf{q}_0 = \mathbf{M}_t \mathbf{Q}_0 \psi = \mathbf{Q}_t \psi. \quad (9)$$

Substituting Eqs. (8) and (9) into Eq. (3) gives the (unregularized) data-driven (DD) approximation of G^{opt} ,

$$G_{\text{DD}}^{\text{opt}} = \max_{\mathbf{q}_0 \in \text{sp}(\mathbf{Q}_0)} \frac{\|\mathbf{q}_t\|^2}{\|\mathbf{q}_0\|^2} = \max_{\psi} \frac{\|\mathbf{Q}_t \psi\|^2}{\|\mathbf{Q}_0 \psi\|^2}. \quad (10)$$

That is, we seek to maximize the energy growth between initial and evolved states within the span of the data, which is equivalent to maximizing the ratio of the energy of the evolved and initial disturbances over all expansion coefficients. Here, the operator $\text{sp}(\cdot)$ returns the span of its matrix argument. Using the definition of the norm in Eq. (4), Eq. (10) may be rewritten as

$$G_{\text{DD}}^{\text{opt}} = \max_{\psi} \frac{\psi^* \mathbf{Q}_t^* \mathbf{W} \mathbf{Q}_t \psi}{\psi^* \mathbf{Q}_0^* \mathbf{W} \mathbf{Q}_0 \psi}. \quad (11)$$

A similar optimization encountered in the context of a data-driven resolvent analysis method of Ref. [24] was solved by formulating a cost function and finding its stationary point. Instead, we solve Eq. (11) by transforming it into a Rayleigh quotient. To this end, we define the Cholesky decomposition

$$\mathbf{Q}_0^* \mathbf{W} \mathbf{Q}_0 = \mathbf{B}^* \mathbf{B}. \quad (12)$$

For now, we assume \mathbf{Q}_0 possesses full column rank, implying that $\mathbf{Q}_0^* \mathbf{W} \mathbf{Q}_0$ is invertible. This assumption is removed with the regularization presented in the next section. Again using $\mathbf{W} = \mathbf{L}^* \mathbf{L}$, Eq. (11) can be expressed as

$$G_{\text{DD}}^{\text{opt}} = \max_{\mathbf{v}} \frac{\mathbf{v}^* \mathbf{B}^{*-1} \mathbf{Q}_t^* \mathbf{L}^* \mathbf{L} \mathbf{Q}_t \mathbf{B}^{-1} \mathbf{v}}{\mathbf{v}^* \mathbf{v}}, \quad (13)$$

where $\mathbf{v} = \mathbf{B}\psi$ (and $\psi = \mathbf{B}^{-1}\mathbf{v}$). The solution to the Rayleigh quotient is

$$G_{\text{DD}}^{\text{opt}} = \sigma_1^2(\mathbf{L}\mathbf{Q}_t\mathbf{B}^{-1}), \quad (14)$$

and this maximum is achieved when $\mathbf{v} = \mathbf{v}_1$, which is the first right singular vector of the matrix $\mathbf{L}\mathbf{Q}_t\mathbf{B}^{-1}$. Finally, the input and output modes can be obtained using the optimal expansion coefficients, $\psi_1 = \mathbf{B}^{-1}\mathbf{v}_1$, as

$$\mathbf{q}_0^{\text{opt}} = \mathbf{Q}_0\psi_1/\|\mathbf{Q}_0\psi_1\|, \quad \mathbf{q}_t^{\text{opt}} = \mathbf{Q}_t\psi_1/\|\mathbf{Q}_t\psi_1\|. \quad (15)$$

In the Appendix, we show that the same result can be obtained using a different formulation inspired by DMD.

B. Regularization

The method presented in the previous section assumes that the data exactly follow Eq. (6). In realistic use cases, the data may be corrupted with noise or contain the influence of nonlinearity; these effects can be modeled as measurement noise and process noise, respectively. This noise can dominate the small eigenvalues in $\mathbf{Q}_0^*\mathbf{W}\mathbf{Q}_0$, and, because $\mathbf{Q}_0^*\mathbf{W}\mathbf{Q}_0$ is in the denominator of the optimization in Eq. (10), these small eigenvalues can dominate the results. Therefore, we introduce a regularization method to improve the robustness of the method by reducing its sensitivity to deviations from the ideal linear dynamics in Eq. (6).

To this end, we regularize the system by adding a constant to the denominator in Eq. (11),

$$G_{\text{DD}}^{\text{opt}} = \max_{\psi} \frac{\psi^* \mathbf{Q}_t^* \mathbf{W} \mathbf{Q}_t \psi}{\psi^* (\mathbf{Q}_0^* \mathbf{W} \mathbf{Q}_0 + \gamma \mathbf{I}) \psi}, \quad (16)$$

where $\gamma \in \mathbb{R}^+$ serves as the regularization parameter, setting a minimum threshold for the eigenvalues of the matrix in the denominator of the optimization. Since smaller eigenvalues are often primarily associated with noise, this regularization method efficiently suppresses such unwanted elements without negatively affecting the larger eigenvalues typically associated with coherent structures within the data. The regularization parameter γ should be selected to fall between the largest and smallest eigenvalues of $\mathbf{Q}_0^*\mathbf{W}\mathbf{Q}_0$ to ensure minimal interference with essential data. Our strategy for choosing γ will be further detailed in Sec. IV. The algorithm then involves performing the Cholesky decomposition

$$(\mathbf{Q}_0^*\mathbf{W}\mathbf{Q}_0 + \gamma\mathbf{I}) = \mathbf{B}^*\mathbf{B}. \quad (17)$$

Since $\mathbf{Q}_0^*\mathbf{W}\mathbf{Q}_0 + \gamma\mathbf{I}$ is now strictly positive definite, the Cholesky decomposition is guaranteed to exist, and \mathbf{B} is guaranteed to be invertible. Compared to the unregularized version, Eq. (17) is the only change in the algorithm. The method, with the regularization, is given in Algorithm 1.

ALGORITHM 1. Data-driven transient growth analysis.

- 1: **Inputs:** $\mathbf{Q}_0, \mathbf{Q}_t, \gamma$
- 2: $\mathbf{B} \leftarrow \text{Chol}(\mathbf{Q}_0^*\mathbf{W}\mathbf{Q}_0 + \gamma\mathbf{I})$
- 3: $[\mathbf{U}, \Sigma, \mathbf{V}] \leftarrow \text{SVD}(\mathbf{L}\mathbf{Q}_t\mathbf{B}^{-1})$
- 4: $\psi_1 \leftarrow \mathbf{B}^{-1}\mathbf{v}_1$
- 5: $G^{\text{opt}} \leftarrow \sigma_1^2$
- 6: $\mathbf{q}_0^{\text{opt}} \leftarrow \mathbf{Q}_0\psi_1/\|\mathbf{Q}_0\psi_1\|$
- 7: $\mathbf{q}_t^{\text{opt}} \leftarrow \mathbf{Q}_t\psi_1/\|\mathbf{Q}_t\psi_1\|$
- 8: **Outputs:** $G^{\text{opt}}, \mathbf{q}_0^{\text{opt}}, \mathbf{q}_t^{\text{opt}}$

Algorithm description. Inputs: \mathbf{Q}_0 , the matrix containing realizations of the initial state; \mathbf{Q}_t , the matrix containing realizations of the states at time t ; γ , the regularization parameter. Outputs: G^{opt} , the optimal energy growth at time t ; $\mathbf{q}_0^{\text{opt}}$ and $\mathbf{q}_t^{\text{opt}}$, the optimal input and output modes for time t .




FIG. 1. One trajectory of the GL equation with $\mu_0 = 0.38$ for a random initial condition. The trajectory displays transient energy growth before an exponential decay.

Assuming fewer snapshots than degrees of freedom ($m < n$), the most computationally costly step in the algorithm is the SVD in line 2, which scales as $O(m^2n)$. This is the same scaling as typical algorithms for POD, so application to large datasets should be straightforward. If the algorithm were run with more snapshots than spatial degrees of freedom ($m > n$), which we do not recommend, then the scaling is determined by the Cholesky decomposition in step 1 and inverse in step 3, both of which scale like $O(m^3)$.

IV. VALIDATION USING A LINEARIZED GINZBURG-LANDAU EQUATION

In this section, we apply our method to data generated by the linearized Ginzburg-Landau (GL) equation to assess our approach in a controlled environment and to determine the best choice for the regularization parameter γ in the presence of noise. The GL equation offers an ideal testbed for this purpose: it supports transient growth akin to the linearized Navier-Stokes equations, and the operator-based theory from Sec. II can be easily applied for comparison against results from the data-driven method from Sec. III.

A. Ginzburg-Landau equation and data matrix construction

The GL equation is

$$\frac{\partial q}{\partial t} = \mathcal{A}q(x) = \left(-v \frac{\partial}{\partial x} + \gamma_g \frac{\partial^2}{\partial x^2} + \mu(x) \right) q, \quad (18)$$

where q is the perturbation about an implied zero base state. The parameters v and γ_g are the flow convection and dissipation coefficients, respectively, and $\mu(x) = \mu_0 - c_u^2 + \mu_2 \frac{x^2}{2}$ produces local exponential growth or decay depending on the sign of $\mu(x)$. Following Ref. [25], we set $v = 2 + 0.4i, \gamma_g = 1 - i, \mu_0 = 0.38, c_u = 0.2$, and $\mu_2 = -0.01$. The equation is spatially discretized using $n = 220$ Hermite polynomials and integrated in time using a fourth-order Runge-Kutta scheme. Figure 1 shows one trajectory of the Ginzburg-Landau system. As time progresses, a transient increase in energy is visible before the eventual asymptotic decay.

To form the data matrices in Eq. (7), we collect m realizations of the data pairs $[\mathbf{q}_0^i, \mathbf{q}_t^i]$, with $t \in [0, T]$, by computing m trajectories of duration T , each starting from a different randomly generated initial condition. We choose initial conditions drawn from a multivariate Gaussian distribution with zero mean and spatial correlation of the form [26]

$$\mathbb{E}[q_0(x_1)q_0^*(x_2)] = \exp\left[-\frac{(x_1 - x_2)^2}{\lambda^2}\right], \quad (19)$$

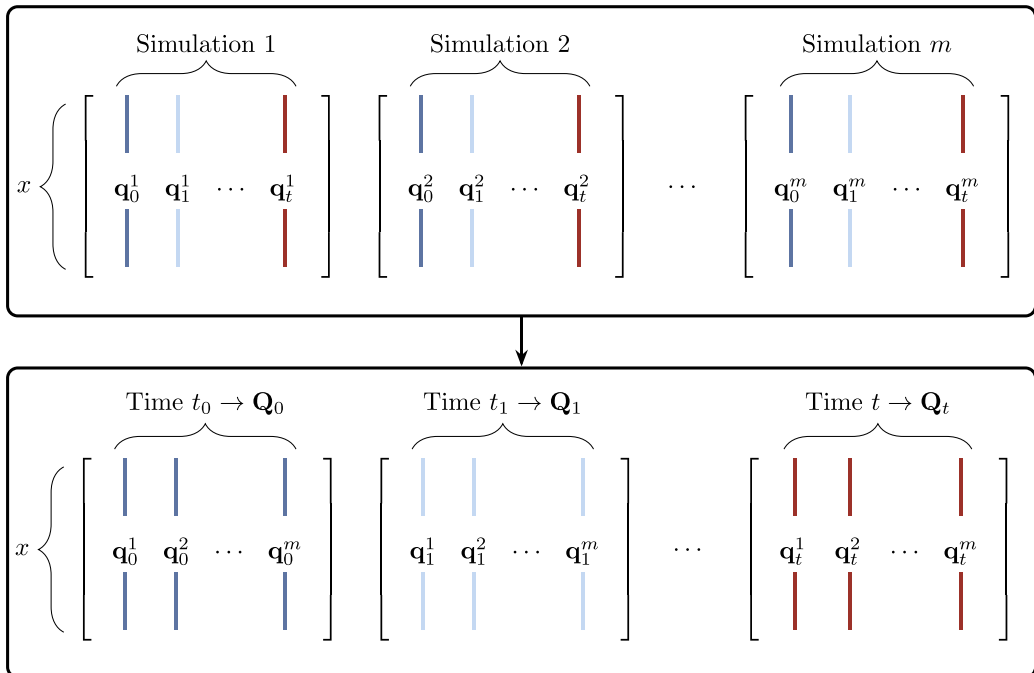


FIG. 2. Constructing the data matrices. The matrices in the first row represent different trajectories of the solution obtained from different initial conditions, with each column color-coded according to its position in time. These columns are sorted by time in the second row, forming \mathbf{Q}_0 and $\mathbf{Q}_t \in \mathbb{C}^{n \times m}$.

with a correlation length of $\lambda = 2$. Each initial condition comprises one column of \mathbf{Q}_0 , and we then evolve each initial condition to generate \mathbf{Q}_t at subsequent times. Alternatively, we could have divided a single, longer trajectory into a series of blocks to obtain the set of realizations used to define the data matrices, since any pair of snapshots still obeys Eq. (6). Figure 2 illustrates this process. The number of realizations should be large enough such that the data matrices are close to spanning the desired optimal input and output modes; this is assessed in Sec. VB.

To assess the resilience of the regularized method to noise, we introduce both measurement and process noise. After spatial discretization, the noisy system is

$$\frac{d\mathbf{q}}{dt} = \mathbf{A}\mathbf{q} + \mathbf{w}, \quad (20)$$

$$\tilde{\mathbf{q}} = \mathbf{q} + \mathbf{n}, \quad (21)$$

where \mathbf{w} and \mathbf{n} represent the process and measurement noise, respectively, and $\tilde{\mathbf{q}}$ denotes the noisy measurement of the state. Both types of noise are spatially Gaussian, as in Eq. (19). The measurement noise is white-in-time, while the process noise is band-limited white noise (between frequencies $\omega = 0.001$ and $\omega = 0.5$) to avoid numerical issues when integrating Eq. (20).

B. Error metrics

We evaluate the data-driven method by comparing its predicted growth and input and output modes against those computed using the standard operator-based approach, which is straightforward to apply to the Ginzburg-Landau problem. We quantify the error using two metrics.

First, the peak growth error is defined as

$$\epsilon_G = \frac{|G_{\max, \text{DD}}^{\text{opt}} - G_{\max}^{\text{opt}}|}{G_{\max}^{\text{opt}}}, \quad (22)$$

where G_{\max}^{opt} indicates the largest optimal growth over all time, and the subscript DD refers to the data-driven results. This metric quantifies the relative discrepancy between the peak transient growth values predicted by the data-driven and operator-based methods.

Second, the error in the input and output modes is defined, respectively, as

$$\epsilon_0^{\text{opt}} = 1 - |\langle \mathbf{q}_{\text{DD}, i}, \mathbf{q}_i \rangle|, \quad \epsilon_r^{\text{opt}} = 1 - |\langle \mathbf{q}_{\text{DD}, o}, \mathbf{q}_o \rangle|, \quad (23)$$

where \mathbf{q} is the state vector. The scalars ϵ_0^{opt} and ϵ_r^{opt} are 1 if the modes are orthogonal and zero if they are the same up to an arbitrary complex phase.

C. Results

In this section, we test our data-driven method using the Ginzburg-Landau equation. First, we look at the effect of the number of realizations, then evaluate the impact of the regularization parameter, and propose a standard for selecting it.

The number of realizations, m , represents the number of independent state vectors evolving over time. The set of realizations should be large enough to accurately span the desired optimal input and output modes. At the same time, the number of realizations may be constrained in practice by the dataset in hand. Moreover, while a larger number of realizations may always appear desirable, we will see that using more data than necessary increases the sensitivity of the method to noise. Overall, these considerations highlight the need for a technique that is effective across a range of m values.

We first evaluate the impact of m in the absence of noise. Typically, one should avoid making m close to or larger than n , as doing so leads to redundancy and computational inefficiency. However, for the sake of exploration, we experimented with $m = [55, 110, 165, 220]$, which ranges from $n/4$ to n . Figure 3 shows the optimal transient growth and its optimal modes for these four values of m using the unregularized method, as well as the operator-based, i.e., ground truth, results. The transient growth plot indicates that the estimation improves as m increases, with smaller m values yielding lower growth values. It is reasonable that with a small m and zero noise, the optimal growth is lower since fewer combinations of vectors are available to maximize over; i.e., the data do not fully span the desired operator-based optimal modes. The output modes are more accurately captured than the input modes. This is a consequence of the growth of the linear operator; the part of each realization of the initial condition that projects onto the optimal input mode is amplified by a factor of G^{opt} , so the mode we are trying to represent is more prevalent in the data. A close match with the input mode is achieved when $n = m$, while the results for the output mode closely align with the analytical outcome for all values of m .

Next, we introduce moderate noise levels ($\mathbb{E}[\|\mathbf{w}\|] = \mathbb{E}[\|\mathbf{n}\|] = 0.03$). Since our initial conditions have expected energy of 1, the expectation of noise can be seen as relative to that of the initial condition. We use a regularization factor of $\gamma = \gamma_0 \lambda_1(\mathbf{Q}_0^* \mathbf{W} \mathbf{Q}_0)$ where $\gamma_0 = [0, 0.01, 0.02, 0.04]$. In other words, we choose the regularization parameter to be proportional to the largest eigenvalue of $\mathbf{Q}_0^* \mathbf{W} \mathbf{Q}_0$ with a constant of proportionality of γ_0 . This approach elucidates the reasonable range of the regularization amplitude: γ_0 should be on the order of a few percent, as smaller values will have little impact and larger values will overwhelm the data. Additionally, this relative definition should aid in making effective values transferable between problems, as shown later.

Figure 4 compares the optimal operator-based and data-driven transient growth for several values of γ_0 and m . Each subfigure shows how the growth changes with m under a particular value of the regularization parameter γ_0 . We observe that larger m increases the potential of the growth regardless of the value of γ_0 , which is due to the availability of more combinations of vectors to maximize over.

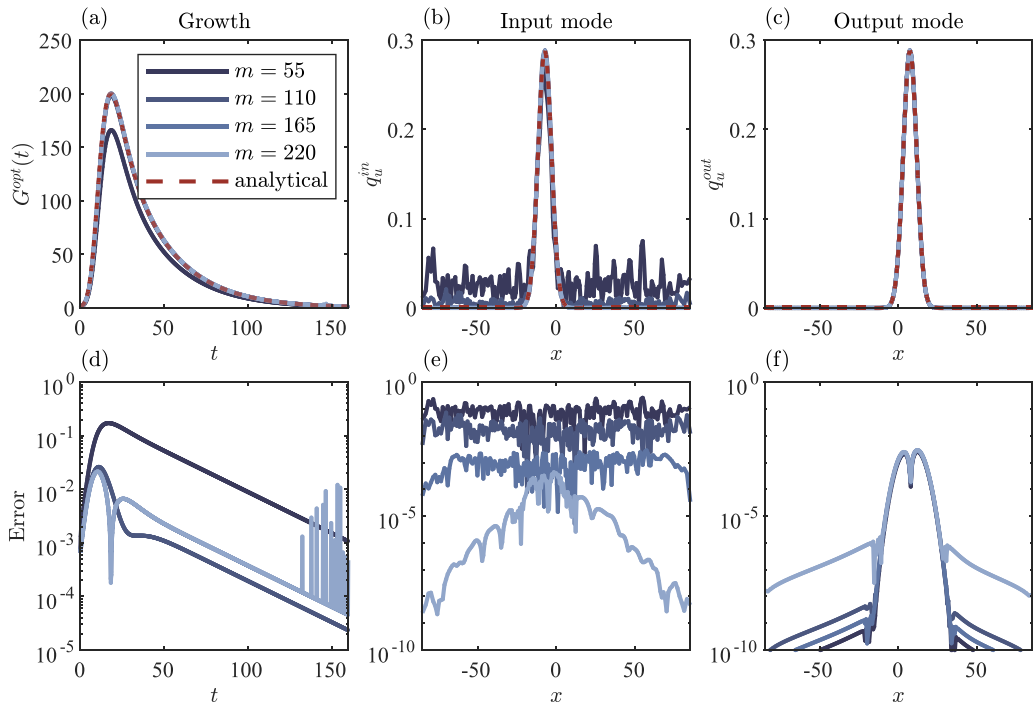


FIG. 3. Comparison between the (unregularized) data-driven and operator-based transient growth results for the Ginzburg-Landau system without noise: (a) optimal growth, (b) optimal input mode, (c) optimal output mode, (d) growth error, (e) input mode error, and (f) output mode error. Error is computed as the difference between the data-driven and operator-based results.

However, a larger m also increases the effect of noise. As shown in Fig. 4, although larger γ_0 helps to suppress the noise, we can still observe that as m increases, the growth results display noisy structures, especially when the growth is small (at $t > 100$).

Figure 5 quantifies the error in the growth, input modes, and output modes for the Ginzburg-Landau equation corrupted by measurement and process noise using the metrics ϵ_G , ϵ_0^{opt} , and ϵ_t^{opt} , respectively, as defined in Eqs. (22) and (23). Unlike the noise-free case, the unregularized method ($\gamma_0 = 0$) does not converge with increasing m ; instead, the error temporarily drops before diverging toward high error levels as more realizations are added. In contrast, the regularized method ($\gamma_0 > 0$) is relatively insensitive to the number of realizations; the growth and input mode errors decrease slightly, and the output mode error increases slightly with increasing m , but the absolute error levels are reasonable in all cases and lowest for the output mode. Additionally, the results are insensitive to the size of the regularization parameter within the range we consider. As in the noise-free case, the error is especially low for the output modes. Overall, these results highlight the critical need for regularization to ensure reliable results when the data are corrupted by noise.

As previously mentioned, we select the regularization parameter γ based on the maximum eigenvalue of $\mathbf{Q}_0^* \mathbf{W} \mathbf{Q}_0$, such that $\gamma = \gamma_0 \lambda_1(\mathbf{Q}_0^* \mathbf{W} \mathbf{Q}_0)$, so that a given value of γ_0 has the same meaning regardless of the number of realizations or the energy of the initial disturbances. The sensitivity of the method to noise stems from the significant disparity between the maximum and minimum eigenvalues of $\mathbf{Q}_0^* \mathbf{W} \mathbf{Q}_0$. The smallest eigenvalues can be easily affected by noise, and, because the algorithm involves inverting the Cholesky factor of $\mathbf{Q}_0^* \mathbf{W} \mathbf{Q}_0$, these eigenvalues can have a significant impact on the outputs of the algorithm.

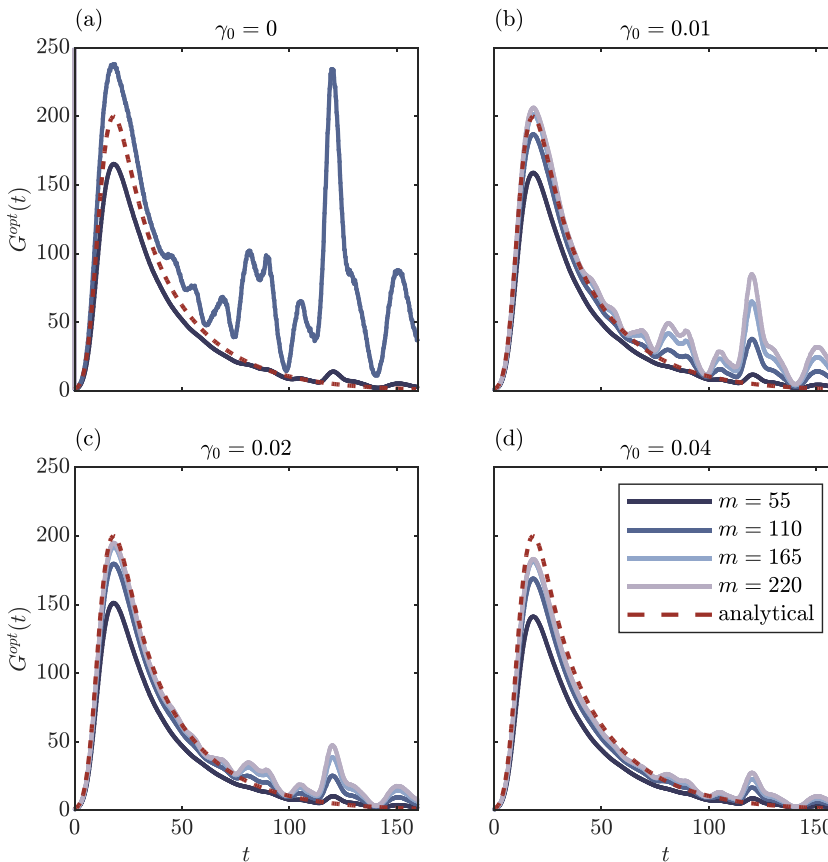


FIG. 4. Data-driven approximations of transient growth for the Ginzburg-Landau system with 3% noise for different numbers of realizations, m , and choices of the regularization parameter γ_0 . The curves for $m = 165$ and $m = 220$ in (a) diverge and are excluded for clarity.

We show the effect of m and γ_0 on our estimates of peak growth and optimal modes in Fig. 6. Figure 6(a) shows how the peak growth error ϵ_G varies with m and γ_0 . We can observe that for $m > 140$ and $0.005 < \gamma_0 < 0.02$, the error becomes mostly independent of m . However, ϵ_0^{opt} and ϵ_t^{opt} exhibit different dependencies on γ and m . As shown in Figs. 6(b) and 6(c), the errors for both input and output modes are nearly independent of γ_0 and m , except when γ_0 is very low (for ϵ_t^{opt}) or when either m or γ_0 is very low (for ϵ_0^{opt}). As previously noted, the input mode is not amplified by the singular values, which explains why ϵ_0^{opt} tends to be higher. In contrast, the output mode can be estimated accurately with fewer realizations, and larger m values negatively affect the results. These observations regarding ϵ_0^{opt} and ϵ_t^{opt} also indicate that there is a flexible range for selecting γ_0 , as long as it is above a specific threshold.

Next, we select a few typical γ_0 values to evaluate their effectiveness across different noise levels. Specifically, in Fig. 7, we show each error metric as a function of the mean amplitude of the measurement and process noise for $\gamma_0 = [0, 0.01, 0.02, 0.04]$. The first column of the figure shows that an unregularized approach ($\gamma_0 = 0$) closely matches analytical predictions in the absence of noise. However, unsurprisingly, its effectiveness decreases as the noise levels increase. The remaining three columns of Fig. 7 show that all three regularized cases ($\gamma_0 > 0$) provide accurate estimates of the maximum transient growth and the optimal modes in the presence of noise. This




FIG. 5. Growth and mode errors for the Ginzburg-Landau system corrupted with 3% measurement and process noise as a function of the number of realizations, m , for several values of the regularization parameter γ_0 : (a) peak growth error ϵ_G , (b) input mode error ϵ_0^{opt} , and (c) output mode error ϵ_t^{opt} .

analysis confirms that, for the GL equation, our method delivers reasonable results within a flexible parameter range.

V. VALIDATION FOR POISEUILLE FLOW

Next, we apply the data-driven method to Poiseuille flow, i.e., the pressure-driven flow between infinite parallel plates. This case provides a physically meaningful benchmark that allows us to address three important considerations. First, recent work [26] showed that, for certain values of the streamwise and spanwise wave numbers, a randomly sampled initial condition is extremely unlikely to achieve near-optimal growth, casting doubt on whether the optimal modes can be obtained from linear combinations of a set of highly suboptimal trajectories. We show that they can be. Second, we show that the data-driven algorithm delivers accurate results for the wide range of $G_{\text{max}}^{\text{opt}}$ values encountered when varying the wave numbers. Third, we show that the regularization parameter γ_0 selected for the previous Ginzburg-Landau problem is also a reasonable choice for Poiseuille flow, suggesting that this default value can be widely applied without case-dependent tuning, including to flows for which no baseline results are readily available, such as the boundary layer considered in Sec. VI.




FIG. 6. Growth and mode errors for the Ginzburg-Landau system corrupted with 3% measurement and process noise as a function of the number of realizations, m , and the regularization parameter γ_0 : (a) peak growth error ϵ_G , (b) input mode error ϵ_0^{opt} , and (c) output mode error ϵ_t^{opt} . In (a), $\epsilon_G = 10^{-2}$ and $\epsilon_G = 10^{-1}$ are marked in bold.




FIG. 7. Growth and mode errors for the Ginzburg-Landau system as a function of the measurement and process noise levels: top row, peak growth error ϵ_G ; middle row, input mode error ϵ_0^{opt} ; and bottom row, output mode error ϵ_t^{opt} . Columns from left to right: $\gamma_0 = [0, 0.01, 0.02, 0.04]$. All cases use $m = 110$ realizations.

A. Problem setup

The velocity-vorticity formulation of the linearized Navier-Stokes equation is [23]

$$\frac{\partial}{\partial t} \begin{bmatrix} \hat{v} \\ \hat{\eta} \end{bmatrix} = -i \begin{bmatrix} \mathcal{L}_{\text{OS}} & 0 \\ \mathcal{L}_{\text{C}} & \mathcal{L}_{\text{SQ}} \end{bmatrix} \begin{bmatrix} \hat{v} \\ \hat{\eta} \end{bmatrix}, \quad (24)$$

where the Orr-Sommerfeld, cross-term, and Squire operators are

$$\mathcal{L}_{\text{OS}} = - \left(\frac{\partial^2}{\partial y^2} - k^2 \right)^{-1} \left[\frac{1}{i\text{Re}} \left(\frac{\partial^2}{\partial y^2} - k^2 \right)^2 - \alpha U \left(\frac{\partial^2}{\partial y^2} - k^2 \right) + \alpha U'' \right], \quad (25)$$

$$\mathcal{L}_{\text{C}} = \beta U', \quad (26)$$

$$\mathcal{L}_{\text{SQ}} = \alpha U - \frac{1}{i\text{Re}} \left(\frac{\partial^2}{\partial y^2} - k^2 \right). \quad (27)$$

Here α and β are streamwise and spanwise wave numbers, respectively, and $k^2 = \alpha^2 + \beta^2$ is the squared wave-number magnitude. The base flow, $U = U(y)$, is only inhomogeneous in the wall-normal direction. The state variables, \hat{v} and $\hat{\eta}$, are the Fourier coefficients of the wall-normal velocity and vorticity perturbations around the base flow U . The notation $(\cdot)'$ represents the wall-normal derivative d/dy , and we use the Chebyshev discretization scheme identical to that in Ref. [5]. We




FIG. 8. Transient growth for Poiseuille flow: operator-based and data-driven optimal growth along with 1000 random trajectories for $\alpha = 0$, $\beta = 2$.

use $N_y = 100$ internal points for the spatial discretization, which results in $n = 200$, since there are two state variables.

We use a kinetic energy norm; in the velocity-vorticity formulation, the kinetic energy is written as

$$e = \frac{1}{4\pi^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{1}{2k^2} \int_{-1}^1 \left\| \frac{\partial}{\partial y} \hat{v} \right\|^2 + k^2 \|\hat{v}\|^2 + \|\hat{\eta}\|^2 dy d\alpha d\beta. \quad (28)$$

The weight matrix in Eq. (4) is then

$$\mathbf{W} = \begin{bmatrix} \mathbf{D}^* \mathbf{W}_c \mathbf{D} + k^2 \mathbf{W}_c & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_c \end{bmatrix}, \quad (29)$$

where \mathbf{D} is the Chebyshev differential operator for computing $\partial/\partial y$, and \mathbf{W}_c is a diagonal matrix that represents the integration weight. The initial condition and the data matrix have the same structure as in Sec. IV A.

B. Poiseuille flow results

Figure 8 shows the trajectories of 1000 random initial conditions, drawn from the same distribution used in Ref. [26], along with the analytical optimal growth and the data-driven approximation thereof for $\text{Re} = 1000$, $\alpha = 0$, and $\beta = 2$. By definition, G^{opt} bounds the trajectories, and, consistent with Ref. [26], all trajectories substantially undershoot this bound. Nevertheless, the data-driven algorithm faithfully predicts the optimal growth curve. This empirically demonstrates that a set of random suboptimal trajectories can span the optimal modes. Mathematically, the key distinction is between any one random sample and the span of many samples: while no single sample is likely to closely match the optimal initial condition, by finding the optimal linear combination of a set of random samples (as the data-driven method does), we can obtain the optimal initial condition within the span of random initial conditions. Figure 9 shows the maximum gain $G_{\text{max}}^{\text{opt}}$ for a range of streamwise and spanwise wave numbers α and β . The data-driven method maintains accuracy over the full range of wave numbers.

Focusing on $\alpha = 0$ and $\beta = 2$, where the maximum transient growth exists, we test how noise and regularization impact the results for Poiseuille flow. Figure 10 shows how our method performs at a moderate level of process noise as a function of the number of realizations, m , and regularization parameter γ_0 . We observe a similar trend to that previously reported for the Ginzburg-Landau equation in Fig. 6. Namely, the data-driven method performs well for a range of nonzero γ_0 values. Similar to Figs. 7 and 11 shows the performance of our method in Poiseuille flow as a function of the measurement and process noise for several values of the regularization parameter γ_0 . We observe a similar trend as shown for the Ginzburg-Landau equation in Fig. 7. Clearly, a




FIG. 9. Maximum optimal growth for Poiseuille flow: (a) operator-based result, (b) data-driven results, and (c) error.

certain level of regularization is essential in the existence of noise. In addition, by defining the level of regularization based on $\lambda_1(\mathbf{Q}_0^* \mathbf{W} \mathbf{Q}_0)$, a relatively consistent strength of regularization is achieved by choosing the same γ_0 . For instance, as shown in Figs. 11(c), 11(g), and 11(k), $\gamma_0 = 0.01$, the same value used for the Ginzburg-Landau problem, produces reasonable and consistent estimations.

VI. SPATIAL TRANSIENT GROWTH IN A TRANSITIONAL BOUNDARY LAYER

The previous section demonstrates the robustness of our method on the GL equation and provides guidance for selecting the regularization parameter γ_0 . In this section, we use this method to study spatial transient growth in a boundary layer. The relevant growth in boundary layers is spatial rather than temporal; i.e., disturbances grow as they advect downstream in the boundary layer, and this process is statistically stationary in time. Thus far, we have discussed temporal transient growth, so we first review spatial transient growth theory in Sec. VIA. We then describe the boundary layer data to which we apply our method in Sec. VIB and describe the data processing required to construct the data matrices in the spatial case in Sec. VIC. Finally, in Sec. VID, we report our results, i.e., the growth and modes for the boundary layer.




FIG. 10. Growth and mode errors for Poiseuille flow corrupted with 5% measurement and process noise as a function of the number of realizations, m , and the regularization parameter γ_0 : (a) peak growth error ϵ_G , (b) input mode error ϵ_0^{opt} , and (c) output mode error ϵ_t^{opt} . In (a), $\epsilon_G = 10^{-2}$ and $\epsilon_G = 10^{-1}$ are marked in bold.




FIG. 11. Growth and mode errors for Poiseuille flow as a function of the measurement and process noise levels: top row, peak growth error ϵ_G ; middle row, input mode error ϵ_0^{opt} ; and bottom row, output mode error ϵ_t^{opt} . Columns from left to right: $\gamma_0 = [0, 10^{-3}, 10^{-2}, 10^{-1}]$. All cases use $m = 150$ realizations.

A. Spatial transient growth

The state vector in the spatial case can be decomposed as

$$\mathbf{q}(x, y, z, t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \hat{\mathbf{q}}(x, y, \beta, \omega) \exp[i(\beta z - \omega t)] d\beta d\omega, \quad (30)$$

where $\hat{\mathbf{q}}(x, y, \beta, \omega)$ is the velocity field Fourier transformed in z and t , β is the spanwise wave number, and ω is the frequency. Upon discretization in the y direction, the state at a certain x , β , and ω is $\hat{\mathbf{q}}(x, \beta, \omega)$. The transformation is standard and is useful because each (β, ω) pair evolves independently; i.e., $\hat{\mathbf{q}}(x, \beta, \omega)$ is a linear function of $\hat{\mathbf{q}}(0, \beta, \omega)$ alone. This is useful for two reasons. First, it reduces computation time, as analyzing the data at a specific (β, ω) pair reduces the original four-dimensional data matrix to a two-dimensional data matrix. Second, it enables examination of flow properties at various frequencies and wave numbers and facilitates comparisons with results from modal analyses.

Analogous to the temporal evolution equation [Eq. (1)], the spatial evolution of the state for each (β, ω) can be written as

$$\frac{d}{dx} \mathbf{q} = \mathbf{A}(x) \mathbf{q}. \quad (31)$$

Here, and in what follows, we have dropped the $(\hat{\cdot})$ and suppressed the arguments for notational convenience; i.e., \mathbf{q} is the y -discretized state at a certain x , β , and ω . An important difference




FIG. 12. Skin-friction coefficient computed from the data compared to laminar and turbulent values for the boundary layer example. The flow begins to transition near $Re_x = 2 \times 10^5$. For this study, we use data to the left of the vertical line.

between the temporal evolution equation in Eq. (1) and the spatial evolution equation in Eq. (31) is that, in the latter, the linear operator is a function of the direction of evolution; i.e., \mathbf{A} is not a function of t in the temporal case but is a function of x in the spatial case. Nevertheless, because the equation is linear, an input-output map exists of the form

$$\mathbf{q}_x = \mathbf{M}_x \mathbf{q}_0. \quad (32)$$

Here, \mathbf{q}_x is the state at downstream position x , \mathbf{q}_0 is the state at an initial upstream position, and \mathbf{M}_x is the linear operator that relates the two. This equation is now entirely analogous to the temporal case in Eq. (2), and we may apply the proposed algorithm. While the spatial propagator is cumbersome to obtain from first principles [13,27], making an operator-based approach difficult, our data-driven approach may be applied with little additional difficulty relative to the temporal case.

B. Boundary layer data

The boundary layer data are obtained from the JHTDB [17–19]. The Reynolds number based on the half-plate thickness L_p is $Re_{L_p} = 800$. The boundary layer undergoes bypass transition beginning at a distance $x \approx 350L_p$ downstream of the leading edge of the plate [19], as indicated by the skin-friction coefficient shown in Fig. 12.

All data in the database are nondimensionalized using the plate half thickness L_p and the origin defined at the leading edge, but these are not the natural choices for stability analyses. In preparation for using our algorithm and comparing our results to those in the stability literature, we transform the data as follows. First, we shift the x coordinate to the best match between the mean streamwise velocity calculated from the data and the Blasius solution [28]. The shifted x coordinate is used to define the Reynolds number,

$$Re_x = \frac{Ux}{\nu} = Re_{L_p} \frac{x}{L_p}. \quad (33)$$

For consistency with the literature, we nondimensionalize the wall-normal and spanwise directions according to the boundary layer displacement thickness from a Blasius profile [20,21]. Specifically, y and z are nondimensionalized as y/δ and z/δ , respectively, where $\delta(x) = \sqrt{\nu x/U_\infty}$ is the x -dependent transverse length scale. The parameter δ is related to the Blasius boundary layer displacement thickness as $\delta^* \approx 1.72\delta$. In the later discussion, δ_{x_0} is the boundary layer thickness

at the first x location in the data, while δ refers to the local thickness at the x location of interest. The dimensionless spanwise wave number is then

$$\beta\delta = (\beta L_p)(\text{Re}_{L_p})^{-1/2} \sqrt{\frac{x}{L_p}}, \quad (34)$$

where βL_p is the spanwise wave number implied by the original nondimensionalization of the database. Finally, the frequency is nondimensionalized as $\omega x_0/U_\infty$.

Transient growth analysis, including our data-driven method, assumes linear dynamics. Accordingly, we confine our analysis to the laminar region of the flow where the nonlinearity is small; i.e., we focus on linear growth of disturbances that preceded transition. This range can be estimated by looking at the skin-friction coefficient given by Ref. [19], as shown in Fig. 12. We used data up to $\text{Re}_x = 2 \times 10^5$, i.e., to the left of the vertical line in the figure.

C. Data matrix construction for the boundary layer data

In this section, we describe how to construct the data matrices required for the spatial data-driven transient growth analysis of the boundary layer. The relevant state vector \mathbf{q} consists of velocity perturbations defined relative to the time- and spanwise-averaged flow, which nearly matches the Blasius solution until the flow transitions to turbulence. There are four main steps: (i) form the state vector, (ii) take a spanwise Fourier transform, (iii) take a temporal Fourier transform, and (iv) form the final data matrix at each streamwise position. These steps are detailed below.

The raw data from the JHTDB is of size $N_v \times N_x \times N_y \times N_z \times N_t$, where $N_v = 3$ is the number of velocity components, $N_x = 556$ is the number of retained streamwise positions, $N_y = 112$ is the number of wall-normal grid points, $N_z = 1024$ is the number of spanwise grid points, and $N_t = 4701$ is the number of temporal snapshots. Since the applicable state vector \mathbf{q} for this problem consists of the three velocity components at each wall-normal grid point, we first reshape the data matrix such that the leading dimension contains the velocity components stacked together, giving a new data matrix of size $n \times N_x \times N_z \times N_t$, where $n = N_v N_y$ is the state dimension. Second, we take the spanwise discrete Fourier transform. Since each spanwise wave number β is decoupled in linear analyses, this reduces the size of the data matrix for each spanwise wave number to $n \times N_x \times N_t$.

Third, we perform the temporal Fourier transforms implied by Eq. (30). Since the flow is ergodic, we partition the single, long time series provided in the database into m shorter blocks, each of which can be interpreted as a realization of the flow. This is analogous to Welch's method for computing power spectra and standard algorithms for spectral proper orthogonal decomposition [29]. Each block contains N_f snapshots, which determines the frequency resolution, i.e., the frequency increment and total number of frequencies, after taking a discrete Fourier transform.

Using blocks of finite length to approximate the infinite Fourier transform implied by Eq. (30) can introduce unwanted effects. In particular, even if the data are entirely linear, Eq. (32) is perfectly satisfied only in the limit of infinite intervals. Physically, this occurs because, if the interval is too short, the initial state \mathbf{q}_0 does not have time to evolve and influence the downstream state \mathbf{q}_x beyond a certain point, leading to decorrelation between the initial and downstream data in each temporal block. This can be mitigated by shifting the temporal position of the blocks used to define the downstream pair of each initial condition. We do so by defining a convection velocity U_c to determine how far into the future to shift a given downstream block to ensure that the data it contains are indeed the downstream response to the initial condition contained within the corresponding input block. Practically, we restrict the possible convection velocities to those values that result in blocks that align with the discrete data, i.e., $U_c = n_x \Delta x / n_t \Delta t$, where n_x and n_t are the numbers of spatial and temporal steps between the input and output blocks and Δx and Δt are the spatial and temporal spacing of the data. Details on determining the optimal U_c can be found in Ref. [28].

Finally, after taking the discrete Fourier transform of each block, the data at a given frequency of interest are extracted from each block and collected into the input and output matrices $\mathbf{Q}_0, \mathbf{Q}_x \in \mathbb{C}^{n \times m}$. The complete procedure described above is schematically depicted in Fig. 13.

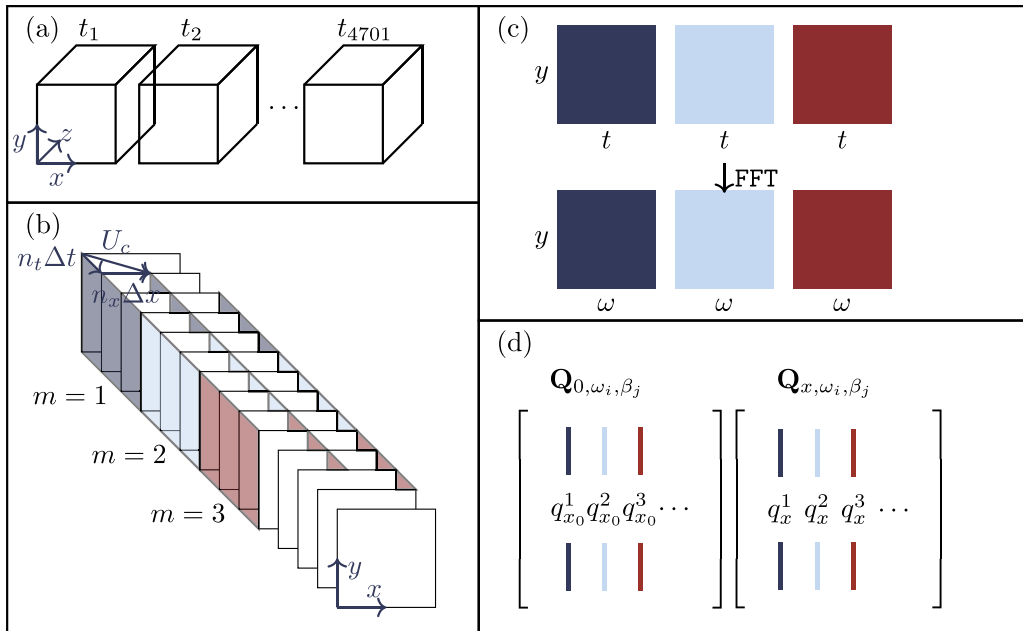


FIG. 13. The procedures for obtaining the input data matrices for our algorithm for the spatially evolving boundary layer. (a) The raw data make up a time series of three-dimensional snapshots. (b) Take a spanwise discrete Fourier transform and select one wave number. The resulting time series is partitioned into a set of realizations; each color indicates a realization. Note that, at downstream x locations, the temporal interval for the same realization occurs later, with the delay determined by the convection velocity U_c to maximize the correlation between the streamwise locations in the realization. (c) Take the temporal Fourier Transform for every realization. (d) Reorganize the output from (c) such that states at the same x , ω , and β are in the same matrix, yielding the inputs to our method.

D. Boundary layer results

Following Refs. [20,21], we focus on $\omega = 0$ and consider a range of β values. We set the regularization parameter to $\gamma_0 = 0.01$, based on the results from the Ginzburg-Landau and Poiseuille flow problems.

The growth over x for three spanwise wave numbers is depicted in Fig. 14. We observe that the growth for $\beta\delta_{x_0} = 0$ and 0.12 is much slower than for $\beta\delta_{x_0} = 0.28$. The slower growth for zero β is consistent with the slow growth of Tollmien-Schlichting waves known to be present at low β . All three curves are similar in magnitude to those given in Refs. [20,21]. It bears mentioning that most curves, even the ones that decay at the end of the interval shown in Fig. 13, grow outside of this interval instead of decaying to zero as $\text{Re}_x \rightarrow \infty$. We believe this can be attributed to the breakdown of the linearity assumption as the flow transitions to turbulence. Also, we note that the magnitude of the growth is sensitive to γ . As we increase or decrease γ by an order of magnitude, the growth also changes by an order of magnitude. Therefore, we rely on our tests with the Ginzburg-Landau equation, which suggests a possible value of γ , though the interpolation of the selection of γ in different cases needs further investigation.

Figure 15 shows the optimal output mode for the three velocity components and the same three $\beta\delta_{x_0}$ values. Except for the $\beta = 0$ results in the first row, we observe that the nonzero β results are consistent over different Re_x and β . The streamwise velocity u and the wall-normal velocity v display a one-peak structure, while the spanwise velocity w has two peaks. At zero β , a two-peak structure is observed for the streamwise velocity disturbance. This structure resembles the




FIG. 14. The spatial transient growth as a function of streamwise location returned by the data-driven method for three spanwise wave numbers.

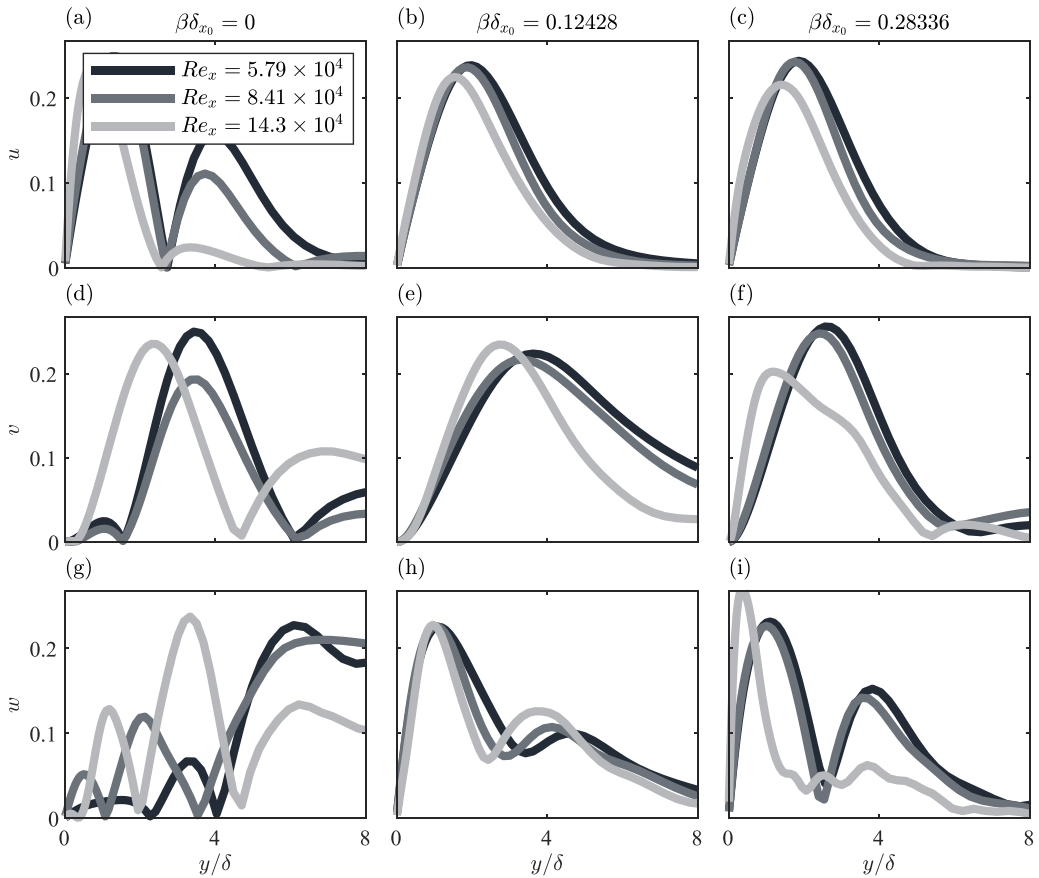


FIG. 15. The output mode for the boundary layer for $\beta\delta_{x_0}$ of 0, 0.12, and 0.28. Each row shows the output modes at different x locations for a single component of the velocity field. A two-peak structure can be observed in (a), while in (b) and (c), there is only one peak.




FIG. 16. Comparison of the streamwise-velocity component of the output mode with the result in Ref. [20] for the boundary layer.

Tollmien-Schlichting (TS) modes described in Ref. [5], and it is well known that modal growth is strongest at $\beta = 0$. References [20,21] have also included the output mode for the streamwise velocity disturbance. The comparison is shown in Fig. 16. The curves for the streamwise velocity have a one-peak structure, which demonstrates good agreement with the results in Refs. [20,21], regardless of x or β . As shown in the figure, the peak of the mode shifts to the right when Re_x increases. The result from Ref. [20] is the limit of $Re_x \rightarrow \infty$, which is consistent with the trend in Fig. 16 that our curve shifts toward their result as Re_x increases.

Figure 17 provides another perspective to investigate the spatial evolution of the output modes. We plot the modes in both streamwise and wall-normal directions for $\beta\delta_{x_0} = 0$ and 0.12. In Figs. 17(b) and 17(c), we can observe that the single-peak structure of the output mode moves away from the wall as the boundary layer thickens. From Fig. 17(a), where we hypothesize modal growth, we observe that the double-peak structure persists for the entire domain, but the top peak significantly reduces in strength. The trough between the two peaks moves away from the wall as we move downstream, while the bottom peak moves closer to the wall. Similar phenomena have been observed in several previous studies [30,31].

In Fig. 18, we show the growth normalized by Re_x against $\beta\delta$; our results are consistent with those of Refs. [20,21] in several ways. These studies showed that G^{opt}/Re_x approaches a constant in the limit $Re_x \rightarrow \infty$ and found that the asymptotic value of G^{opt}/Re_x is maximized at $\beta\delta = 0.45$. If we ignore the growth at $\beta\delta = 0$ in Fig. 18, we can observe that there exists an optimal growth peak




FIG. 17. Spatial evolution of output mode for (a) $\beta\delta_{x_0} = 0$, (b) $\beta\delta_{x_0} = 0.12$, and (c) $\beta\delta_{x_0} = 0.28$.

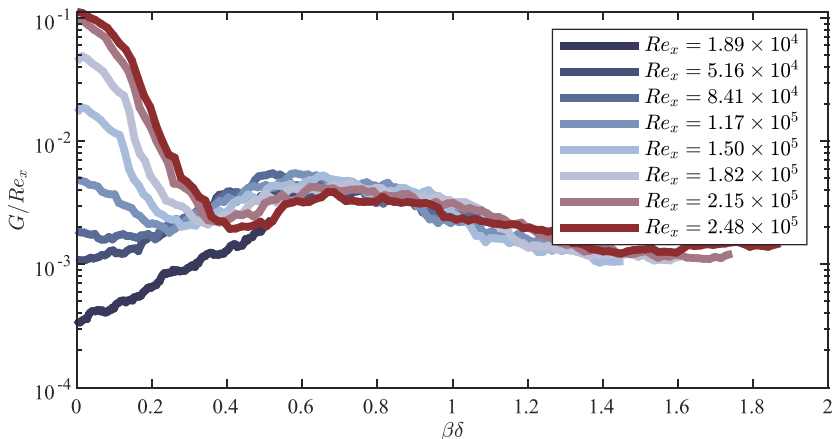


FIG. 18. Growth as a function of the spanwise wave numbers at different Re_x and $\omega = 0$. The curves are smoothed with `movmean` in MATLAB. The optimal growth for various Re_x has a local peak at about $\beta\delta = 0.52$. The large growth at $\beta\delta = 0$ reflects the substantial modal growth possible at low β .

at $\beta\delta \approx 0.52$. We believe that the growth near $\beta\delta = 0$ for the higher Re_x values is due to modal growth, consistent with the differing mode shapes observed in Fig. 15. Additionally, the curves roughly overlap within the range of Re_x in the figure. What is more, the shape and magnitude of the growth are also similar to findings in Refs. [20,21]. However, direct comparisons are not possible due to the difference in problem setup.

Finally, Fig. 19 shows the growth over different frequencies, which has a maximum growth at $\omega = 0$ for all values of Re_x . This is the same as stated in Ref. [21] and justifies focusing on the result with $\omega = 0$.

VII. LIMITATIONS AND OUTLOOKS

The data-driven framework proposed in this paper is designed to replicate traditional, operator-based linear transient growth analysis. As such, a fundamental assumption of our method is that the underlying dynamics of the data are linear or near linear. This limits the method's applica-




FIG. 19. Growth as a function of the frequency for different Re_x . The maximum growth is observed at $\omega = 0$.

bility to two primary scenarios: (i) data generated by a linear solver, as in our Ginzburg-Landau example, where the method provides a valuable matrix-free and adjoint-free tool for analysis; and (ii) data from nonlinear simulations or experiments representing small perturbations about a laminar equilibrium, as in our transitional-boundary-layer example such that their evolution is governed primarily by the linearized dynamics. In the first scenario, the data-driven method is valuable in that it provides a matrix-free and adjoint-free algorithm for transient growth analysis. The second scenario is practically relevant because it includes transitional flows, the class of flows to which transient growth analysis is most often applied.

Working with data that include the impact of nonlinearity poses a challenge. Our regularization scheme is not designed to explicitly separate linear and nonlinear dynamics. Instead, it operates under the assumption that small deviations from linearity can be treated as a form of process noise, which the regularization can partially mitigate. The effectiveness of this approach was demonstrated with the forced Ginzburg-Landau equation and Poiseuille flow. However, in the boundary layer data, where strong freestream disturbances trigger bypass transition, the flow is inherently nonlinear. This approximation is a potential source of the sensitivity we observe between the regularization parameter and the calculated optimal growth. Nevertheless, even in this near-linear regime, our method successfully identifies the structure of the optimal output modes, indicating its potential to extract dominant linear mechanisms from complex data.

It is worth noting that this linearity assumption is common in the foundational literature on related data-driven techniques, such as data-driven resolvent analysis [32]. We argue that this assumption is particularly well suited to transient growth analysis, which is most often applied to transitional flows where linear mechanisms can play a dominant role, in contrast to resolvent analysis, which is frequently applied to fully turbulent flows where nonlinearity is pervasive.

A promising avenue for future research is the development of methods that can relax the strict linearity assumption. This would broaden the applicability of data-driven transient growth analysis to more strongly nonlinear phenomena. Such an approach might involve incorporating techniques from nonlinear system identification or adapting recent efforts in the data-driven resolvent community that aim to distinguish between linear and nonlinear dynamics. Developing such an algorithm is a significant undertaking and lies beyond the scope of the present work.

VIII. CONCLUSIONS

We have developed a data-driven algorithm for optimal transient growth analysis. The appeal of the proposed method is twofold. First, the linearized Navier-Stokes operator required in traditional transient growth calculations may be cumbersome to obtain from a large code or may not be readily available, as in the case of an experiment. Similarly, matrix-free methods require adjoint capabilities. Conversely, the data required to apply the proposed method are readily available in many cases, making the method easier to use in many scenarios. Second, the proposed method scales much more favorably to large problems. While the traditional approach scales cubically in the spatial dimension due to the matrix exponential and SVD, the data-driven approach scales linearly in the spatial dimension and quadratically in the number of snapshots used—the same scaling as obtaining POD modes, which is not difficult for large problems.

We validated the ability of the method to produce accurate estimates of these quantities for a linearized Ginzburg-Landau system, where the growth and modes can be computed analytically. We also used the Ginzburg-Landau system to determine a reasonable choice for the regularization parameter in the method. We then applied our method to the transitional boundary layer to identify spatial transient growth. Our prediction for the output mode of the streamwise velocity component aligns closely with operator-based results. From the output modes, we observed a clear distinction between the two-peak structure for modal growth and the one-peak structure for transient growth. We also identified a spanwise wave number that contributes significantly to transient energy growth. The discrepancies between our results and the literature include differences in the growth over the streamwise location and the spanwise wave number that maximizes energy growth. We believe that

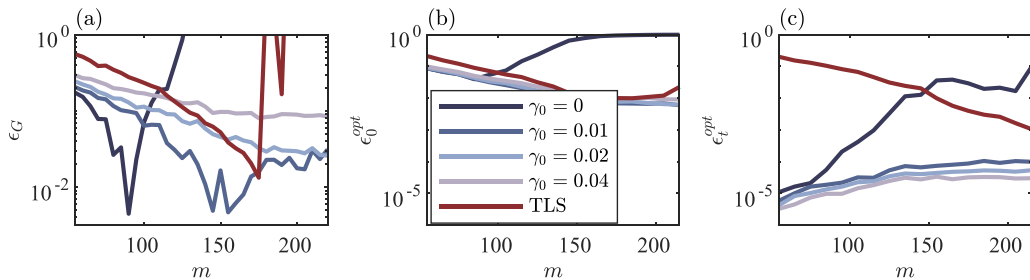


FIG. 20. Same as Fig. 5, with the additional red curve representing the error of TLS.

much of the discrepancy is caused by the strong freestream disturbances triggering bypass transition present in the data, and that we would have seen better performance had we chosen a case in which linear growth played a more prominent role in the transition process.

One particularly promising application of the method is studying hypersonic boundary layer transition. The transition process for high-speed flows is sensitive to numerous factors, and it can be challenging to include all of the relevant physics within a linear stability solver that could be used to conduct an operator-based transient growth analysis. In contrast, excellent data exist for these flows that could be used as input to our data-driven transient-growth algorithm.

ACKNOWLEDGMENT

This work was supported in part by ONR Grant No. N00014-24-1-2033.

DATA AVAILABILITY

The data that support the findings of this article are openly available [18].

APPENDIX: LEAST-SQUARES APPROXIMATION METHOD

Here, we show another formulation of the data-driven method inspired by dynamic mode decomposition (DMD) [33] that is equivalent to the formulation of the method presented in the main text in the case of no regularization. In addition to offering another perspective on the method that may be more intuitive to some readers, this formulation suggests alternative regularization strategies.

The assumption of linearity guarantees that there is a matrix \mathbf{M} such that

$$\mathbf{Q}_t = \mathbf{M}\mathbf{Q}_0. \quad (\text{A1})$$

For illustration purposes, we show the unweighted case here, i.e., $\mathbf{W} = \mathbf{I}$. Analogous to DMD, we approximate \mathbf{M} as

$$\hat{\mathbf{M}} = \mathbf{Q}_t\mathbf{Q}_0^+ \approx \mathbf{M}, \quad (\text{A2})$$

where $\mathbf{Q}_0^+ = (\mathbf{Q}_0^*\mathbf{Q}_0)^{-1}\mathbf{Q}_0^*$ (as we have assumed \mathbf{Q}_0 has full column rank). The optimal growth is then approximated using this estimation of the evolution operator as

$$G_{\text{LSA}}^{\text{opt}}(t) = \sigma_1^2(\hat{\mathbf{M}}). \quad (\text{A3})$$

The similarity between this method and DMD lies in the way they both approximate the linear operator relating two data matrices; the differences between the two methods are that (i) the corresponding columns in the data matrices in DMD are a single time step apart, whereas here they are separated by the time between the input and output disturbance, and (ii) DMD takes an eigendecomposition of the resulting linear operator and we take an SVD.




FIG. 21. Peak growth error ϵ_G for the Ginzburg-Landau system as a function of measurement and process noise levels: (a) unregularized, (b) the regularization from Sec. III B with $\gamma_0 = 0.02$, and (c) TLS regularization.

This formulation is equivalent to that presented in the main text in the unregularized case, and this can be shown as follows. Leveraging the assumption that \mathbf{Q}_0 has full column rank, we first take the Cholesky decomposition

$$\mathbf{B}_{\text{UW}}^* \mathbf{B}_{\text{UW}} = \mathbf{Q}_0^* \mathbf{Q}_0, \quad (\text{A4})$$

where the subscript is a reminder that we are in the unweighted case. Now, Eq. (A3) may be rewritten as

$$G_{\text{LSA}}^{\text{opt}}(t) = \sigma_1^2(\mathbf{Q}_t \mathbf{B}_{\text{UW}}^{-1} \mathbf{B}_{\text{UW}}^{*-1} \mathbf{Q}_0^*). \quad (\text{A5})$$

The square of the largest singular value of any matrix is the largest eigenvalue of the same matrix multiplied by its Hermitian conjugate. Using this fact, we have

$$G_{\text{LSA}}^{\text{opt}}(t) = \lambda_1(\mathbf{Q}_t \mathbf{B}_{\text{UW}}^{-1} \mathbf{B}_{\text{UW}}^{*-1} \mathbf{Q}_0^* \mathbf{Q}_0 \mathbf{B}_{\text{UW}}^* \mathbf{B}_{\text{UW}}^{-1} \mathbf{Q}_t^*). \quad (\text{A6})$$

Using Eq. (A4), this simplifies to

$$G_{\text{LSA}}^{\text{opt}}(t) = \lambda_1(\mathbf{Q}_t \mathbf{B}_{\text{UW}}^{-1} \mathbf{B}_{\text{UW}}^{*-1} \mathbf{Q}_t^*). \quad (\text{A7})$$

Finally, using the correspondence between singular values and eigenvalues, we have

$$G_{\text{LSA}}^{\text{opt}}(t) = \sigma_1^2(\mathbf{Q}_t \mathbf{B}_{\text{UW}}^{-1}). \quad (\text{A8})$$

Equation (A8) is equivalent to Eq. (14) in the case of no weight matrix, i.e., $\mathbf{L} = \mathbf{I}$. It can also be shown that

$$\mathbf{q}_0^{\text{opt}} = \mathbf{Q}_0 \psi_{\text{UW},1} / \|\mathbf{Q}_0 \psi_{\text{UW},1}\|, \quad \mathbf{q}_t^{\text{opt}} = \mathbf{Q}_t \psi_{\text{UW},1} / \|\mathbf{Q}_t \psi_{\text{UW},1}\|, \quad (\text{A9})$$

where $\psi_{\text{UW},1} = \mathbf{B}_{\text{UW}}^{-1} v_{\text{UW},1}$, and $v_{\text{UW},1}$ is the first column of the right singular vector of $\mathbf{Q}_t \mathbf{B}_{\text{UW}}^{-1}$. Additionally, $\mathbf{q}_t^{\text{opt}}$ is equivalent to the left singular vector of $\mathbf{Q}_t \mathbf{B}_{\text{UW}}^{-1}$, as shown in Ref. [28].

The DMD-based approach in the weighted case approximates \mathbf{M} by first transforming to a coordinate system where the 2-norm is equivalent to the \mathbf{W} norm in the original coordinates, then approximating \mathbf{M} with a pseudoinverse in these coordinates. Mathematically,

$$\hat{\mathbf{M}} = \mathbf{Y}_t \mathbf{Y}_0^+, \quad (\text{A10})$$

where $\mathbf{Y}_t = \mathbf{L} \mathbf{Q}_t$ and $\mathbf{Y}_0 = \mathbf{L} \mathbf{Q}_0$ are the matrices of weighted states. The growth is again given by the SVD of $\hat{\mathbf{M}}$, and this may be shown to be equivalent to

$$G_{\text{LSA}}^{\text{opt}}(t) = \sigma_1^2(\mathbf{L} \mathbf{Q}_t \mathbf{B}_{\text{UW}}^{-1}), \quad (\text{A11})$$

which is the same as the growth in the unregularized, weighted case in the main text. The modes can also be shown to be the same.




FIG. 22. The output mode for the boundary layer at various streamwise locations and for three spanwise wave numbers computed using the TLS method.

The potential advantage of this equivalence and the DMD-based approach is that we can now regularize it with methods originally designed for DMD. Here, we choose the total least squares (TLS) method from Ref. [34]. As shown in Figs. 20 and 21, this method generally works for the test case with the Ginzburg-Landau equation. Figure 20 is equivalent to Fig. 5 in the main text but with the addition of the results for TLS regularization. Note that the $\gamma_0 = 0$ case is equivalent to the DMD-based method without regularization. Compared with the regularization method in previous sections, the total least squares method depends only on the number of realizations, m , but produces a less accurate result in general. As shown in Fig. 20, more realizations allow the total least squares method to get better estimations of the modes. However, the growth estimated by the total least squares is very sensitive to noise when m is large. Additionally, the TLS regularization produces less accurate output modes than the method in the main text for all m values. Figure 21 plots the peak growth error, ϵ_G , under different measurement and process noise with $m = 110$. In this figure, we can observe that the TLS method can produce reasonable results over the range of noise. Although the ϵ_G for TLS is higher than the regularization proposed in Sec. III B, it still outperforms the unregularized method.

For the boundary layer data, the TLS regularization can only capture the optimal output modes as shown in Fig. 22. The output mode results are consistent with the other method. However, the growth result in Fig. 23 is obviously dominated by noise. The cause of this is not entirely clear; one potential explanation could be the insufficient number of realizations, m . Since the TLS method involves initial projection into a low-dimensional POD space, this step could significantly reduce the




FIG. 23. The spatial transient growth for the boundary layer as a function of streamwise location for three spanwise wave numbers computed using the TLS method.

accuracy of the result. In addition, treating the nonlinearities as noise may exceed the capabilities of the TLS method, since its development in Ref. [34] aims to suppress process noise without considering nonlinearities.

-
- [1] H. Bénard, Les tourbillons cellulaires dans une nappe liquide. Méthodes optiques d'observation et d'enregistrement, *J. Phys. Theor. Appl.* **10**, 254 (1901).
 - [2] L. Rayleigh, LIX. On convection currents in a horizontal layer of fluid, when the higher temperature is on the under side, *London, Edinburgh, Dublin Philos. Mag. J. Sci.* **32**, 529 (1916).
 - [3] G. I. Taylor, VIII. Stability of a viscous liquid contained between two rotating cylinders, *Philos. Trans. R. Soc. London, Ser. A* **223**, 289 (1923).
 - [4] N. Tillmark and P. H. Alfredsson, Experiments on transition in plane Couette flow, *J. Fluid Mech.* **235**, 89 (1992).
 - [5] P. J. Schmid, D. S. Henningson, and D. Jankowski, Stability and transition in shear flows. Applied mathematical sciences, Vol. 142, *Appl. Mech. Rev.* **55**, B57 (2002).
 - [6] P. G. Drazin and W. H. Reid, *Hydrodynamic Stability* (Cambridge University Press, Cambridge, UK, 2004), Chap. 8, pp. 124–250.
 - [7] L. Böberg and U. Brösa, Onset of turbulence in a pipe, *Z. Naturforsch. A* **43**, 697 (1988).
 - [8] K. M. Butler and B. F. Farrell, Three-dimensional optimal perturbations in viscous shear flow, *Phys. Fluids* **4**, 1637 (1992).
 - [9] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll, Hydrodynamic stability without eigenvalues, *Science* **261**, 578 (1993).
 - [10] E. Reshotko, Transient growth: A factor in bypass transition, *Phys. Fluids* **13**, 1067 (2001).
 - [11] M. F. De Pando, D. Sipp, and P. J. Schmid, Efficient evaluation of the direct and adjoint linearized dynamics from compressible flow solvers, *J. Comput. Phys.* **231**, 7739 (2012).
 - [12] A. Towne, R. Bhagwat, Y. Zhou, J. Jung, E. Martini, P. Jordan, D. Audiffred, I. Maia, and A. Cavalieri, Resolvent-based estimation of wavepackets in turbulent jets, *30th AIAA/CEAS Aeroacoustics Conference (2024)* (American Institute of Aeronautics and Astronautics/Council of European Aerospace Societies, Reston, VA, 2024), p. 3413.
 - [13] A. Towne and T. Colonius, One-way spatial integration of hyperbolic equations, *J. Comput. Phys.* **300**, 844 (2015).
 - [14] M. Zhu and A. Towne, Recursive one-way Navier-Stokes equations with PSE-like cost, *J. Comput. Phys.* **473**, 111744 (2023).
 - [15] L. Sirovich, Turbulence and the dynamics of coherent structures. I. Coherent structures, *Q. Appl. Math.* **45**, 561 (1987).
 - [16] A. Dotto, D. Barsi, D. Lengani, and D. Simoni, A data-driven optimal disturbance procedure for free-stream turbulence induced transition, *Phys. Fluids* **34**, 124108 (2022).
 - [17] E. Perlman, R. Burns, Y. Li, and C. Meneveau, Data exploration of turbulence simulations using a database cluster, in *Proceedings of the 2007 ACM/IEEE Conference on Supercomputing* (ACM, New York, 2007), pp. 1–11.
 - [18] Y. Li, E. Perlman, M. Wan, Y. Yang, C. Meneveau, R. Burns, S. Chen, A. Szalay, and G. E. and, A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence, *J. Turbul.* **9**, N31 (2008); data available at <http://turbulence.pha.jhu.edu>.
 - [19] T. A. Zaki, From streaks to spots and on to turbulence: Exploring the dynamics of boundary layer transition, *Flow Turbul. Combust.* **91**, 451 (2013).
 - [20] P. Andersson, M. Berggren, and D. S. Henningson, Optimal disturbances and bypass transition in boundary layers, *Phys. Fluids* **11**, 134 (1999).
 - [21] P. Luchini, Reynolds-number-independent instability of the boundary layer over a flat surface: Optimal perturbations, *J. Fluid Mech.* **404**, 289 (2000).

- [22] B.-T. Chu, On the energy transfer to small disturbances in fluid flow (Part I), *Acta Mech.* **1**, 215 (1965).
- [23] S. C. Reddy and D. S. Henningson, Energy growth in viscous channel flows, *J. Fluid Mech.* **252**, 209 (1993).
- [24] A. S. Towne, Advancements in jet turbulence and noise modeling: Accurate one-way solutions and empirical evaluation of the nonlinear forcing of wavepackets, Ph.D. thesis, California Institute of Technology, 2016.
- [25] S. Bagheri, D. S. Henningson, J. Høpfner, and P. J. Schmid, Input-output analysis and control design applied to a linear model of spatially developing flows, *Appl. Mech. Rev.* **62**, 020803 (2009).
- [26] P. Frame and A. Towne, Beyond optimal disturbances: A statistical framework for transient growth, *J. Fluid Mech.* **983**, A2 (2024).
- [27] A. Towne, G. Rigas, O. Kamal, E. Pickering, and T. Colonius, Efficient global resolvent analysis via the one-way Navier–Stokes equations, *J. Fluid Mech.* **948**, A9 (2022).
- [28] Z. Kai, Data-driven transient growth analysis: With application to a transitional boundary layer, Master’s thesis, University of Michigan, 2024.
- [29] A. Towne, O. T. Schmidt, and T. Colonius, Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis, *J. Fluid Mech.* **847**, 821 (2018).
- [30] J. H. M. Fransson, L. Brandt, A. Talamelli, and C. Cossu, Experimental study of the stabilization of Tollmien–Schlichting waves by finite amplitude streaks, *Phys. Fluids* **17**, 054110 (2005).
- [31] S. Westerbeek, J. A. Franco Sumariva, T. Michelis, S. Hein, and M. Kotsonis, Linear and nonlinear stability analysis of a three-dimensional boundary layer over a hump, in *2023 AIAA Science and Technology Forum and Exposition* (American Institute of Aeronautics and Astronautics, National Harbor, Maryland, 2023), p. 0678.
- [32] B. Herrmann, P. J. Baddoo, R. Semaan, S. L. Brunton, and B. J. McKeon, Data-driven resolvent analysis, *J. Fluid Mech.* **918**, A10 (2021).
- [33] P. J. Schmid, Dynamic mode decomposition of numerical and experimental data, *J. Fluid Mech.* **656**, 5 (2010).
- [34] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz, On dynamic mode decomposition: Theory and applications, *J. Comput. Dyn.* **1**, 391 (2014).