

Accurate and Efficient Autonomic Closure for Turbulent Flows

by

Abhinav Kshitij

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved April 2019 by the
Graduate Supervisory Committee:

Werner J.A. Dahm, Chair
Peter E. Hamlington
Marcus Herrmann
Jeonglae Kim
Yulia T. Peet

ARIZONA STATE UNIVERSITY

May 2019

© 2019 Abhinav Kshitij

All Rights Reserved

ABSTRACT

Autonomic closure is a new general methodology for subgrid closures in large eddy simulations that circumvents the need to specify fixed closure models, and instead allows a fully-adaptive self-optimizing closure. The closure is autonomic in the sense that the simulation itself determines the optimal relation at each point and time between any subgrid term and the variables in the simulation, through the solution of a local system identification problem. It is based on highly generalized representations of subgrid terms having degrees of freedom that are determined dynamically at each point and time in the simulation. This can be regarded as a very high-dimensional generalization of the dynamic approach used with some traditional prescribed closure models, or as a type of “data-driven” turbulence closure in which machine-learning methods are used with internal training data obtained at a test-filter scale at each point and time in the simulation to discover the local closure representation.

In this study, *a priori* tests were performed to develop accurate and efficient implementations of autonomic closure based on particular generalized representations and parameters associated with the local system identification of the turbulence state. These included the relative number of training points and bounding box size, which impact computational cost and generalizability of coefficients in the representation from the test scale to the LES scale. The focus was on studying impacts of these factors on the resulting accuracy and efficiency of autonomic closure for the subgrid stress. Particular attention was paid to the associated subgrid production field, including its structural features in which large forward and backward energy transfer are concentrated.

More than five orders of magnitude reduction in computational cost of autonomic closure was achieved in this study with essentially no loss of accuracy, primarily by using efficient frame-invariant forms for generalized representations that greatly reduce the number of degrees of freedom. The recommended form is a 28-coefficient representation that provides subgrid stress and production fields that are far more accurate in terms of structure and statistics than are traditional prescribed closure models.

TABLE OF CONTENTS

	Page
LIST OF TABLES	v
LIST OF FIGURES	vi
CHAPTER	
NOMENCLATURE	xi
1 INTRODUCTION	1
1.1 Filtered Equations and Resulting Subgrid Terms	2
1.2 Filtered Transport Equations and Their Subgrid Terms	2
1.3 Example: The Subgrid Stress	3
1.4 Subgrid Stress Closures and LES Energetics	5
1.5 Forward/Backward Scatter in $P(\mathbf{x}, t)$ and LES “Blowup”	6
1.6 A New Approach: Autonomic Closure	9
1.7 Present Study	10
1.7.1 Objectives	11
1.7.2 Organization of the Dissertation	12
2 THE AUTONOMIC CLOSURE METHODOLOGY	14
2.1 Traditional Prescribed Closure Models vs. Autonomic Closure	14
2.2 Generalized Representations at the Test and LES Scales	14
2.3 Formulation of the Optimization Problem	17
2.4 Implementation Choices in Autonomic Closure	19
2.5 Anticipated Effects of Implementation Choices	20
3 METRICS TO ASSESS ACCURACY OF AUTONOMIC CLOSURE	22
3.1 Pseudo-LES Fields, Test Fields, and Resulting Stress Fields	23
3.2 Statistical Comparisons of $\tau_{ij}(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$	24
3.3 Support Fields for the Subgrid Production $P(\mathbf{x}, t)$	25
3.4 Subgrid Production Support-Density Fields $G(\mathbf{x}, t)$	26

CHAPTER	Page
3.5 Support-Density Metrics \mathcal{M}_1 and \mathcal{M}_2	27
3.6 Computational Time Scaling	28
4 PERFORMANCE OF SERIES-BASED REPRESENTATIONS	31
4.1 Computational Times	32
4.2 Local vs. Nonlocal Representations	32
4.3 Velocity-Pressure vs. Velocity-Only Representations	56
4.4 First-Order vs. Second-Order Representations	58
4.5 Effect of Number of Training Points and Their Separation	61
4.6 Average Subgrid Dissipation Rates	62
4.7 Results from Highly-Sheared Homogeneous Turbulence	63
4.8 Recommended Series-Based Representation	69
4.9 Comparison with Traditional Closure Models	70
5 INVARIANT REPRESENTATIONS IN AUTONOMIC CLOSURE	81
5.1 Frame-Invariant Representation Theory	81
5.2 Invariant Representations of τ_{ij} in the Stencil Velocities \mathbf{u}_m	82
5.3 Invariant Representations of τ_{ij} in \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$, and $\nabla\mathbf{R}$	86
5.3.1 Invariance-Preserving Tensor Representations of τ_{ij}	87
5.3.1.1 Linear representations in \mathbf{S}	87
5.3.1.2 Nonlinear representations in \mathbf{S} and \mathbf{R}	88
5.3.2 Toward the Most-General Frame-Invariant Representation	90
5.3.3 Frame-Invariant Combinations of \mathbf{I} , \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$	92
5.3.3.1 Complete set of rank-2 tensor polynomial bases	92
5.3.3.2 The tensor basis in (5.13) and (5.14) is minimal	95
5.3.3.3 Tensor elements \mathbf{A}_k in terms of \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$	96
5.3.3.4 Symmetric and antisymmetric tensors \mathbf{M}_k and \mathbf{W}_k	99
5.3.3.5 Symmetric rank-2 tensor polynomial basis $\mathbf{m}_S^{(i)}$	101

CHAPTER	Page
5.3.3.6 Incompressible case: Deviatoric basis $\mathbf{m}_S^{(i)}$	103
5.3.4 Second-Order Truncation of the Complete Tensor Basis	104
6 PERFORMANCE OF INVARIANT REPRESENTATIONS	107
6.1 Frame-Invariant Representations for F_{ij}	107
6.2 Error-Free Differentiation and Multiplication	109
6.3 Specific Implementations of Invariant Representations	110
6.4 Resulting Computational Times	112
6.5 Resulting Performance of Frame-Invariant Representations	112
6.6 Eigenvalue Distributions in the Generalized Representations	117
6.7 Recommended Frame-Invariant Representation	124
7 CONCLUSIONS AND IMPLICATIONS FOR LES	125
7.1 Summary	125
7.2 Major Conclusions from This Study	126
REFERENCES	129

LIST OF TABLES

Table	Page
<p>1 Implementations of Autonomic Closure Considered in Chapter 4, Showing Case Number and Code, Number N of Degrees-of-freedom in F_{ij}, Relative Bounding Box Size N^3, Number M of Training Points in Bounding Box, Number of Training Points Per Degree-of-freedom M/N, Relative Training Point Spacing, Computational Time, and Ratio of Autonomic and True Average Subgrid Dissipation. Case Codes: First Character N: Non-colocated, C: Collocated; Second Character L: Local, G: Nonlocal; Third Character 1: First-order, 2: Second-order; Primes: Velocity-pressure Cases; DS = Dynamic Smagorinsky Model, BD = Bardina Scale Similarity Model; + From (4.1), Actual Times 10-20X Longer Due to Memory Management for Large Matrix Operations.</p>	30
<p>2 Frame-invariant Implementations of Autonomic Closure Considered in Chapter 5, Showing Code, Number N of Degrees-of-freedom in F_{ij}, Relative Bounding Box Size N^3, Number M of Training Points in Bounding Box, Relative Training Point Spacing, Number M' of Training Data in Bounding Box, Number of Training Data per Degree-of-freedom M'/N, and Computational Time</p>	111

LIST OF FIGURES

Figure	Page	
1	Typical <i>a Priori</i> Test of Autonomic Closure, Showing (a) True Subgrid Stress Field $\tau_{ij}(\mathbf{x}, t)$ and (c) True Subgrid Production Field $P(\mathbf{x}, t)$ Compared to Corresponding Results from Implementation of Autonomic Closure [59] for (b) Subgrid Stress Field $\tau_{ij}^{\mathcal{F}}(\mathbf{x}, t)$ and (D) Subgrid Production Field $P^{\mathcal{F}}(\mathbf{x}, t)$	5
2	Stencils $\widehat{\mathbf{S}}$ and $\widetilde{\mathbf{S}}$ Each Centered on Point-of-interest \mathbf{x} (Red Dot); (a) Test-grid Stencil $\widehat{\mathbf{S}}$ on Which \mathbf{h}_{ij} Is Determined, (b) Les-grid Stencil $\widetilde{\mathbf{S}}$ on Which Resulting \mathbf{h}_{ij} Is Used to Evaluate $\tau_{ij}^{\mathcal{F}}(\mathbf{x})$	16
3	Typical Subgrid Production Field $P(\mathbf{x}, t)$, Showing (a) True Field, (b) Result from Autonomic Closure, (c) True Support-density Field $G(\mathbf{x}, t)$, and (d) Corresponding Support-density Field from Autonomic Closure.	24
4	Subgrid Production Fields $P(\mathbf{x}, t)$ (Leftmost Column) and Associated Support-density Fields $G(\mathbf{x}, t)$ Filtered at Successive Scale Ratios $\Delta_{\Gamma}/\widetilde{\Delta}$, Showing (Top Row) True Production Field and (Bottom Row) Result from Autonomic Closure. . .	25
5	Scaling of Computational Time T for Autonomic Closure with Number N of Degrees of Freedom in F_{ij} and Number M of Training Points in $\widehat{\mathbf{V}}$ in (2.4) and (2.5), Showing Theoretical Scaling Vs. Actual Computational Time for Cases in Table.1 .	33
6	Local Vs Nonlocal Implementations. Typical Comparison of True Fields (Left) and Autonomic Closure, Case 5a (Right), Showing (a) Normal Subgrid Stresses, (b) Shear Subgrid Stress, and (c) Subgrid Production $P(\mathbf{X}, T)$	37
7	Local Vs Nonlocal Implementations. Typical Comparison of True Fields (Left) and Autonomic Closure, Case 6a (Right), Showing (a) Normal Subgrid Stresses, (b) Shear Subgrid Stress, and (c) Subgrid Production $P(\mathbf{x}, t)$	40
8	Velocity-pressure Vs Velocity-only Implementations. Typical Comparison of True Fields (Left) and Autonomic Closure, Case 4a (Right), Showing (a) Normal Subgrid Stresses, (b) Shear Subgrid Stress, and (c) Subgrid Production $P(\mathbf{x}, t)$	43

Figure	Page
9 Velocity-pressure Vs Velocity-only Implementations. Typical Comparison of True Fields (Left) and Autonomic Closure, Case 5b (Right), Showing (a) Normal Subgrid Stresses, (b) Shear Subgrid Stress, and (c) Subgrid Production $P(\mathbf{x}, t)$	46
10 First-order Vs Second-order Implementations. Typical Comparison of True Fields (Left) and Autonomic Closure, Case 1a (Right), Showing (a) Normal Subgrid Stresses, (b) Shear Subgrid Stress, and (c) Subgrid Production $P(\mathbf{x}, t)$	49
11 First-order Vs Second-order Implementations. Typical Comparison of True Fields (Left) and Autonomic Closure, Case 1b (Right), Showing (a) Normal Subgrid Stresses, (b) Shear Subgrid Stress, and (c) Subgrid Production $P(\mathbf{x}, t)$	52
12 Typical Comparison of (Left Column) True Subgrid Stress Component Field $\tau_{ij}(\mathbf{x}, t)$ and Subgrid Production Field $P(\mathbf{x}, t)$ with (Right Column) Results from Recommended Implementation (Case 3a) of Autonomic Closure; Accuracy Is Comparable to Figure 1 but at Nearly 500x Lower Computational Cost.	55
13 Comparison of Nonlocal (Blue) and Local (Red) Implementations of Autonomic Closure; (Top) Pdfs of (Left) Typical Subgrid Stress Component τ_{12} and (Right) Subgrid Production P Versus (Black) Exact Results, and (Bottom) \mathcal{M}_1 and \mathcal{M}_2 Variations with Scale Ratio $\Delta_\Gamma/\tilde{\Delta}$ for Subgrid Production Field.	56
14 Comparison of Velocity-pressure (Blue) Vs. Velocity-only (Red) Implementations of Autonomic Closure; (Top) Pdfs of (Left) Typical Subgrid Stress Component τ_{12} and (Right) Subgrid Production P Versus (Black) Exact Results, and (Bottom) \mathcal{M}_1 and \mathcal{M}_2 Variation in (11a,b) with Scale Ratio $\Delta_\Gamma/\tilde{\Delta}$	58
15 Comparison of Second-order (Blue) Vs. First-order (Red) Implementations of Autonomic Closure; (Top) Pdfs of (Left) Typical Subgrid Stress Component τ_{12} and (Right) Subgrid Production P Versus (Black) Exact Results, and (Bottom) \mathcal{M}_1 and \mathcal{M}_2 Variation in (11a,b) with Scale Ratio $\Delta_\Gamma/\tilde{\Delta}$	59

Figure	Page
16 Effect of Number of Training Points per Degree of Freedom in Implementations of Autonomic Closure with $M/N \leq 1$ (Blue) and $M/N \gg 1$ (Red); (Top) Pdfs of (Left) Typical Subgrid Stress Component τ_{12} and (Right) Subgrid Production P Versus (Black) Exact Results, and (Bottom) \mathcal{M}_1 and \mathcal{M}_2 Variation in (11a,b) with Scale Ratio $\Delta_\Gamma/\tilde{\Delta}$ for Subgrid Production Support-density Fields.	60
17 Two-dimensional Anisotropy State, $\text{IIIa} = 0.2686$ And $-\text{IIa} = 0.7876$, (Marked In Black) Characterize Highly Anisotropic Flows ($S^* = 8.8$), Typically Found In The Log-layer ($y^+ \sim 100$) Of Turbulent Boundary Layers At Moderately High Re.	64
18 Results for Highly Sheared Homogeneous Turbulence, Showing Typical Comparisons of True Fields (Left) and Autonomic Closure (Case 3a) (Right) for (a) Normal Stresses, (b) Shear Stress, and (c) Subgrid Production $P(\mathbf{x}, t)$	68
19 Results for Highly-sheared Turbulence, Showing Comparisons of Pdfs from True Production Fields (Black) with Production Fields from Autonomic Closure (Red) and from the Dynamic Smagorinsky Model (Blue).	68
20 \mathcal{M}_1 (Left) and \mathcal{M}_2 (Right) Metrics for the Production Fields from the HIT and the HST Cases Are Near-identical, Indicating That the Autonomic Closure Retains Structural Similarity Broadly Across the Entire Spectrum of Turbulent Flows.	69
21 Comparison of Probability Densities for Subgrid Stress Component Fields $\tau_{ij}(\mathbf{x}, t)$ from <i>a Priori</i> Tests of Autonomic Closure, Showing Resulting Distributions for True Subgrid Stress Fields (Black) and from Recommended Implementation (Case 3a) of Autonomic Closure (Red).	71
22 Comparison of Probability Densities for Subgrid Production Field $P(\mathbf{x}, t)$ from <i>a Priori</i> Tests of Autonomic Closure, Showing Results for True Subgrid Production Fields (Black) and from Recommended Implementation (Case 3a) of Autonomic Closure (Red).	72

Figure	Page
23 Typical Comparison of Autonomic Closure and Traditional Prescribed Closure Models, Showing (Figure 23a to Figure 23g) Typical Subgrid Stress $\tau_{ij}(\mathbf{x}, t)$ and (Figure 23b) Subgrid Production $P(\mathbf{x}, t)$; (a) True Fields, (b) Results from Autonomic Closure (Case 3a), (c) Results from Dynamic Smagorinsky (DS) Model, and (d) Results from Bardina Scale-similarity (BD) Model; All Are for Same Scale Ratio $\alpha \equiv \Delta_\Gamma/\tilde{\Delta} = 2$	77
24 Typical Comparison of Autonomic Closure and Traditional Prescribed Closure Models, Showing (Top) Pdfs of Typical Subgrid Stress Component τ_{Ij} and Subgrid Production P , and (Bottom) M_1 and M_2 Variation with Scale Ratio $\delta_\gamma/\tilde{\delta}$ for Subgrid Production Support-density Fields; True Fields (Black), Autonomic Closure (Blue), Dynamic Smagorinsky Model (Red Dashed), Bardina Scale-similarity Model (Red Dotted); All Are for Same Scale Ratio $\alpha \equiv \delta_\gamma/\tilde{\delta} = 2$	78
25 Scaled Timings of Frame-invariant Implementations	112
26 Comparison of (a) True Stress Fields τ_{11} and Those Computed by (b) CL24, (c) Smith-UN, (d) Smith-UC, (e) Smith-20, (f) Pope-11, (g) ASU-8, and (h) Smith-5. . .	114
27 Comparison of (a) True Stress Fields τ_{12} and Those Computed by (b) CL24, (c) Smith-UN, (d) Smith-UC, (e) Smith-20, (f) Pope-11, (g) ASU-8, and (h) Smith-5. . .	115
28 Comparison of (a) True Production Fields and Those Computed by (b) CL24, (c) Smith-UN, (d) Smith-UC, (e) Smith-20, (f) Pope-11, (g) ASU-8, and (h) Smith-5. . .	116
29 Comparison of Pdfs for (a) Stress Field τ_{12} , and (b) Production Field P	117
30 Inverse System for CL24.	120
31 Inverse System for SmithUC.	120
32 Inverse System for SmithUN.	121
33 Inverse System for Smith5.	121
34 Inverse System for ASU8.	122
35 Inverse System for Pope11.	122

Figure	Page
36 Inverse System for Smith19.	123

NOMENCLATURE

SYMBOL	DEFINITION
G	Support-density function
\mathbf{k}	Wave vector
L_x, L_y, L_z	Grid Lengths
\mathbf{M}, \mathbf{m}	Symmetric tensor
M	Number of training points
\mathcal{M}_1	Correlation metric
\mathcal{M}_2	Normalized <i>rms</i> difference
N	Degrees of freedom
n	Bounding box size
P	Subgrid production rate
Re_λ	Taylor scale Reynolds number
\mathbf{R}	Rotation rate tensor
\mathbf{S}	Strain rate tensor
S^*	Nondimensional shear parameter, $S^* = Sk/\epsilon$
T, t	Computational time
T_{ij}	Test-scale stress
V_B	Bounding box volume
v_{rms}	RMS velocity fluctuation
\mathbf{W}	Anti-symmetric tensor
x, y, z	Cartesian coordinate system
y^+	Dimensionless wall distance
II, III	Reynolds stress tensor invariants

α	Scale ratio
Γ	Convolution filter kernel
γ	Threshold fraction
Δ	Grid scale
∇	Gradient operator
$\partial u_i / \partial x_j$	Velocity gradient
$\partial S_{ij} / \partial x_k$	Strain rate gradient
$\partial R_{ij} / \partial x_k$	Rotation rate gradient
ϵ	Average subgrid dissipation rate
λ	Damping coefficient
Σ	Support function threshold
τ_{ij}	Subgrid stress
φ	Subgrid scalar variable

Chapter 1

INTRODUCTION

Large eddy simulation (LES) is being increasingly applied to complex flows [1–11] as the increasing availability of computing power and its decreasing cost make the computational burden of LES more acceptable. At the same time, there have been technical advances in the underlying methodology, such as modern wall treatments [12–15], that are further reducing the computational cost of LES to acceptable levels. These developments are making multiphysics large eddy simulations of complex flows increasingly practical [1], in which the simulations address not only the underlying turbulent flow but also include numerous other coupled physical processes, such as transport of conserved scalars [16–19], droplet and particle dynamics [20–23], phase changes [22, 23], reacting species [24–26], heat transfer [24–26], and other phenomena.

Each physical process introduces governing equations, such as equations for conservation of mass, momentum, energy, and scalars, expressed as a combination of linear and nonlinear terms in the velocity field $\mathbf{u}(\mathbf{x}, t)$ and various scalar fields $\varphi(\mathbf{x}, t)$. Due to the spatial filtering inherent in LES, each nonlinear term in these equations creates an associated subgrid term when the equations are written in the corresponding resolved fields $\tilde{\mathbf{u}}(\mathbf{x}, t)$ and $\tilde{\varphi}(\mathbf{x}, t)$. Each of these subgrid terms must be related to the resolved fields to obtain a closed set of equations. Closure of subgrid terms has traditionally been done by means of prescribed subgrid models [27–33] that typically involve substantial *ad hoc* treatments. Errors introduced by these models can be important contributors to the overall error in results from large eddy simulations [10, 34–38]. For this reason, developing a general method that provides accurate subgrid closures for large eddy simulations has been a central focus area of turbulence research over the past several decades.

1.1 Filtered Equations and Resulting Subgrid Terms

In general, any governing equation can be written as a sum of linear terms $L(\mathbf{u}, \varphi)$ and nonlinear terms $N(\mathbf{u}, \varphi)$ as

$$L(\mathbf{u}, \varphi) + N(\mathbf{u}, \varphi) = 0 \quad (1.1)$$

Applying a suitable spatial filter $(\widetilde{\cdot})$ [29, 39–41] having characteristic length scale $\widetilde{\Delta}$ then gives the corresponding governing equation in the resolved fields $\widetilde{\mathbf{u}}(\mathbf{x}, t)$ and $\widetilde{\varphi}(\mathbf{x}, t)$ as

$$L(\widetilde{\mathbf{u}}, \widetilde{\varphi}) + N(\widetilde{\mathbf{u}}, \widetilde{\varphi}) = -[\widetilde{N(\mathbf{u}, \varphi)} - N(\widetilde{\mathbf{u}}, \widetilde{\varphi})], \quad (1.2)$$

where the right side in (1.2) are subgrid terms that appear because the linearity of $L(\mathbf{u}, \varphi)$ allows $\widetilde{L(\mathbf{u}, \varphi)} = L(\widetilde{\mathbf{u}}, \widetilde{\varphi})$ while the nonlinearity of $N(\mathbf{u}, \varphi)$ leads to $\widetilde{N(\mathbf{u}, \varphi)} \neq N(\widetilde{\mathbf{u}}, \widetilde{\varphi})$. All such subgrid terms must be dealt with in a way that provides a closed set of governing equations in the resolved variables $\widetilde{\mathbf{u}}$ and $\widetilde{\varphi}$. To date, such closure has been achieved by introducing prescribed subgrid models based on various approximations that relate subgrid terms to parameters that are obtainable from the resolved variables. Many such prescribed subgrid models have been proposed for subgrid terms in LES [27–34, 42–57]. Errors from these prescribed models, as revealed for instance in *a priori* tests, can be substantial even for subgrid terms that are fundamental to LES, such as the subgrid stress [27–30, 52–54].

1.2 Filtered Transport Equations and Their Subgrid Terms

For example, of particular relevance for multi-physics simulations involving fluid flows are governing transport equations for the conservation of mass, momentum, and various conserved scalars φ . For clarity taking the density and fluid transport properties to be constant, the resulting filtered forms of these transport equations can respectively be written as

$$\frac{\partial \widetilde{u}_i}{\partial x_i} = 0 \quad (1.3a)$$

$$\frac{\partial \widetilde{u}_i}{\partial t} + \frac{\partial}{\partial x_j} \widetilde{u}_i \widetilde{u}_j = -\frac{1}{\rho} \frac{\partial \widetilde{p}}{\partial x_i} + \nu \frac{\partial^2 \widetilde{u}_i}{\partial x_j \partial x_j} - \frac{\partial}{\partial x_j} [\widetilde{u_i u_j} - \widetilde{u}_i \widetilde{u}_j] \quad (1.3b)$$

$$\frac{\partial \widetilde{\varphi}}{\partial t} + \frac{\partial}{\partial x_j} \widetilde{u}_j \widetilde{\varphi} = D_\varphi \frac{\partial^2 \widetilde{\varphi}}{\partial x_j \partial x_j} - \frac{\partial}{\partial x_j} [\widetilde{u_i \varphi} - \widetilde{u}_i \widetilde{\varphi}] \quad (1.3c)$$

where the subgrid terms are indicated by square brackets.

Since the mass conservation equation is strictly linear, the effect of the filter produces no subgrid terms in the resulting filtered mass conservation equation in (1.3a). However, the non-linear advection term in both the momentum and scalar transport equations leads to a subgrid term in each corresponding filtered transport equation in (1.3b,c). In the momentum equation the resulting subgrid term is called the *subgrid stress tensor* – it accounts for the momentum exchange between the resolved and subgrid scales in the simulation. In the scalar transport equation the resulting subgrid term is called the *subgrid scalar flux vector* – it accounts for scalar transport between the resolved and subgrid scales in the simulation. The modeling of these and similar subgrid terms that arise in a wide range of multi-physics large eddy simulations has been a focus of fluid dynamics research for at least the past 40 years.

1.3 Example: The Subgrid Stress

Taking the subgrid stress as an example, in the original momentum equation the nonlinear product $u_i u_j$ in the advection term $\partial(u_i u_j)/\partial x_j$ leads via (1.2) to the subgrid stress in the resolved-scale momentum equation of the form

$$[\widetilde{N(\mathbf{u})} - N(\widetilde{\mathbf{u}})] = \widetilde{u_i u_j} - \widetilde{u_i} \widetilde{u_j} \equiv \tau_{ij}. \quad (1.4)$$

Widely used models for the subgrid stress τ_{ij} include the basic Smagorinsky model [42], the dynamic Smagorinsky model [27, 46–49], the scale-similarity model [27, 43–45], and mixed models that combine a scale similarity model with a dissipative model [50–57]. All of these produce substantial errors in their representation of $\tau_{ij}(\mathbf{x}, t)$, as has been shown in *a priori* tests [10, 27–30, 34]. The accuracy with which any such subgrid model represents $\tau_{ij}(\mathbf{x}, t)$ from the resolved variables $\widetilde{\mathbf{u}}(\mathbf{x}, t)$ and $\widetilde{p}(\mathbf{x}, t)$ determines how accurately it accounts for the detailed space- and time-varying momentum exchange and associated kinetic energy exchange between the resolved and subgrid scales in a simulation.

If simulating the flow field were the only objective, then continued reliance on such traditional prescribed subgrid stress models might be acceptable, since the filter scale $\widetilde{\Delta}$ could sim-

ply be made sufficiently small (albeit at greater computational cost) so that errors introduced by inaccuracies from the $\tau_{ij}(\mathbf{x}, t)$ model would not substantially affect much larger scales of the flow. However, LES is increasingly being used to simulate not only the flow field $\tilde{\mathbf{u}}(\mathbf{x}, t)$, but also other physical processes occurring in the flow, many of which depend strongly on the smallest scales in the resolved flow field, such as diffusion-limited chemical reactions [2–5, 24–26] and droplet/particle transport and agglomeration [6, 7, 20–23]. In such cases, errors introduced at the smallest resolved scales from a substantially inaccurate τ_{ij} model can create large errors throughout the resolved fields of primary interest. Achieving high fidelity in such multiphysics simulations may therefore require new approaches for representing subgrid terms, including the subgrid stress, that are substantially more accurate than current prescribed subgrid modeling approaches.

It will be shown here that a recently proposed alternative approach [58, 59] to subgrid closure, referred to as “autonomic closure”, can be implemented in computationally efficient ways to enable representation of subgrid fields with significantly greater accuracy across all resolved scales than is possible with traditional prescribed subgrid closure models. Figure 1 shows typical results from *a priori* tests of an implementation [59] of autonomic closure, comparing true subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$ and associated subgrid kinetic energy production fields $P(\mathbf{x}, t)$ to the corresponding results from autonomic closure. The implementation of autonomic closure in Figure 1, while undeniably accurate, is however far too computationally costly for practical use. This dissertation describes autonomic closure in detail and identifies specific implementations that retain comparable accuracy in $\tau_{ij}(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$ as seen in Figure 1, even near the smallest resolved scales, but unlike the implementation in Ref. [59] are computationally efficient enough for practical use in large eddy simulations.




Figure 1. Typical *a priori* test of autonomic closure, showing (a) true subgrid stress field $\tau_{ij}(\mathbf{x}, t)$ and (c) true subgrid production field $P(\mathbf{x}, t)$ compared to corresponding results from implementation of autonomic closure [59] for (b) subgrid stress field $\tau_{ij}^{\mathcal{F}}(\mathbf{x}, t)$ and (d) subgrid production field $P^{\mathcal{F}}(\mathbf{x}, t)$.

1.4 Subgrid Stress Closures and LES Energetics

With regard to resolved-scale energetics and computational stability, even more important than the subgrid stress itself is the corresponding subgrid kinetic energy production field

$$P(\mathbf{x}, t) = -\tau_{ij}\tilde{S}_{ij}, \quad (1.5)$$

where \tilde{S}_{ij} is the resolved strain rate tensor, since this determines both the accuracy of energy exchange between the resolved and subgrid scales and the computational stability of the simulation itself. It is known from *a priori* tests [10, 27–30, 52–54, 60] that true $P(\mathbf{x}, t)$ fields in turbulent flows are highly intermittent, consisting of widely varying values that can be locally positive or negative, with magnitudes of P in highly concentrated regions far exceeding the

true average subgrid dissipation rate $\langle P(\mathbf{x}, t) \rangle \equiv \epsilon$. Large positive P values concentrated in these regions correspond to local instantaneous kinetic energy transfer from the resolved scales into the subgrid scales (“forward scatter”), while negative values give the local rate of energy transfer from subgrid scales into the resolved scales (“backscatter”).

Such large magnitudes of forward and backward scatter in $P(\mathbf{x}, t)$ can be seen for example in Figure 1c, which shows the strong spatial intermittency that is characteristic of subgrid production fields. Large positive (red) and large negative (blue) P values are clustered in relatively compact regions that occupy a small fraction of the domain in which the most intense forward and backward scatter are concentrated. For a $\tau_{ij}(\mathbf{x}, t)$ closure to accurately represent the precise space- and time-varying exchange of momentum and energy between resolved and subgrid scales, including near the smallest scales, it must allow forward and backward scatter in $P(\mathbf{x}, t)$ while providing the correct statistical distribution of P values and accurately representing the highly intermittent regions in which large positive and negative $P(\mathbf{x}, t)$ values are concentrated.

1.5 Forward/Backward Scatter in $P(\mathbf{x}, t)$ and LES “Blowup”

Some τ_{ij} models that allow for backscatter can induce instability in a simulation if their resulting $P(\mathbf{x}, t)$ fields are insufficiently accurate [27, 29, 40, 50–57]. This can occur if the average subgrid dissipation rate $\langle P(\mathbf{x}, t) \rangle$ is too low relative to the true average rate ϵ of energy transfer into the subgrid scales. For this reason, scale-similarity models and other models are often combined with a purely dissipative model to give a sufficiently large average subgrid dissipation rate to maintain computational stability. However, even when the average subgrid dissipation rate is sufficiently large, a subgrid model could still induce instability if it produces incorrectly large local values of backscatter, or if the regions in which large values of backscatter are concentrated occur in the wrong locations or at the wrong times, or persist for too long. At the same time, the $\tau_{ij}(\mathbf{x}, t)$ closure must also produce the correct values of forward scatter in the correct locations at the correct times and for the correct durations. Remarkably little is

known, even today, about the precise dynamics that lead to sudden local exponential “blowup” in the kinetic energy at one or more points in a large eddy simulation.

Subgrid models that are purely dissipative, such as the basic Smagorinsky model [42], ensure stability but are unable to accurately represent the detailed momentum and energy exchange between resolved and subgrid scales in a simulation [27–30, 42]. The dynamic Smagorinsky model [27, 46–49] and various scale similarity models [27, 43–45] include backscatter to increase simulation fidelity, especially near the smallest resolved scales. However, some of these models produce insufficiently large $\langle P(\mathbf{x}, t) \rangle$ to maintain computational stability, and others may lead to instability if the modeled backscatter is too strong or appears at the wrong locations or the wrong times [35–37, 48]. For this reason, some of these models introduce backscatter limiters and other *ad hoc* adjustments to increase subgrid dissipation in order to ensure stable simulations.

Yet it is tautological that if a τ_{ij} closure exactly produces the complete details of the true $\tau_{ij}(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$ fields in *a priori* tests, then in the absence of numerical errors from the LES code [35–38] the closure will be stable despite the required large backscatter. Presumably a τ_{ij} closure can be less than perfect in this respect and still maintain stability. However, beyond the requirement that $\langle P(\mathbf{x}, t) \rangle = \epsilon$, relatively little is known about how accurately the subgrid stress $\tau_{ij}(\mathbf{x}, t)$ must be represented to avoid backscatter instability while providing high fidelity in the detailed momentum and energy transfer even near the smallest scales of a simulation.

Although backscatter may be needed to achieve simulation accuracy in all resolved scales, the presence of backscatter alone is meaningless unless the backscattered energy is introduced in about the right places and the right times, and at the right magnitudes and for the right durations. Due to the highly intermittent nature of $P(\mathbf{x}, t)$ fields, as seen in Figure 1c, it may not be possible in an *a priori* sense to have exact point-by-point agreement between the true subgrid production field and that resulting from a closure for the subgrid stress. However, the subgrid production field from a subgrid stress closure should nevertheless be structurally similar to the true subgrid production field, having large values of forward and backward scatter concentrated

in regions at the same locations and of the same size and shape as in the true $P(\mathbf{x}, t)$ field. With the ergodic hypothesis this also ensures that the closure will produce concentrations of forward and backward scatter at the correct times and for the correct durations.

Based on these considerations, it is expected that accuracy across all resolved scales can be achieved while maintaining computational stability if the following three conditions are met:

1. The τ_{ij} closure should produce $P(\mathbf{x}, t)$ fields that, in *a priori* tests, provide a sufficiently large average subgrid dissipation rate, namely $\langle P(\mathbf{x}, t) \rangle \geq \epsilon$, and should ideally give $\langle P(\mathbf{x}, t) \rangle = \epsilon$.
2. Resulting $P(\mathbf{x}, t)$ fields from the τ_{ij} closure should produce similar statistical distributions of positive and negative values (forward and backward scatter) as do the true $P(\mathbf{x}, t)$ fields.
3. $P(\mathbf{x}, t)$ fields from the τ_{ij} closure should be structurally similar to the true $P(\mathbf{x}, t)$ fields in *a priori* tests, with large magnitudes of forward and backward scatter concentrated in regions at the right spatial locations and of the right size and shape, despite the highly intermittent nature of $P(\mathbf{x}, t)$ preventing the two fields from being exactly identical on a point-by-point basis.

Although *a priori* tests of $P(\mathbf{x}, t)$ alone cannot determine if a closure will provide stable simulations, such tests are the most direct way to assess the accuracy at all resolved scales in the subgrid stress fields τ_{ij} and the associated subgrid production fields from a given closure. Implementing a closure in an LES code for *a posteriori* tests introduces additional effects from the code that can obscure insights into the underlying accuracy of the subgrid closure [27]. For these reasons, *a priori* tests are used here to assess the accuracy of various implementations of autonomic closure in representing τ_{ij} and $P(\mathbf{x}, t)$ fields. For implementations that are found to be accurate in these tests, subsequent *a posteriori* tests can be conducted to determine their stability when implemented in an LES code.

Particular attention is paid here not only to the resulting statistical distributions of forward and backward scatter in $P(\mathbf{x}, t)$, but also to the detailed spatial structure of regions in which large magnitudes of forward and backward scatter are concentrated. Results in the following sections show that efficient implementations of autonomic closure can represent momentum and energy exchange between resolved and subgrid scales, across essentially all resolved scales, far more accurately than do traditional prescribed subgrid closure models.

1.6 A New Approach: Autonomic Closure

An entirely different approach to subgrid closures, termed “autonomic closure”, was recently proposed [58, 59] to circumvent the need to specify a particular fixed parametric closure relation, and instead allow a fully-adaptive self-optimizing closure methodology. The closure is autonomic in the sense that the simulation itself determines the optimal relation at each point and time between any subgrid term and the primitive variables in the simulation, through the solution of a local system identification problem. The closure can be nonparametric in the sense that a generalized representation for the subgrid term is formulated in the resolved primitive variables of the simulation, rather than in parameters formed from them that are presumed to be appropriate, or it can be parametric in the sense that the generalized representation is in terms of various parameters that can be formed from the primitive variables. In both cases, the resulting large number of degrees of freedom in the generalized representation allows autonomic closure to freely adapt to widely varying local turbulence conditions via high-dimensional system identification of the local degree of nonlinearity, nonlocality, nonequilibrium, and other characteristics [61, 62] of the turbulence state at each point and time in the simulation.

Autonomic closure can be regarded as a high-dimensional generalization of the dynamic approach used with various traditional prescribed closure models [27, 46, 47]. Viewed another way it can be regarded as a type of “data-driven” turbulence closure [63–71], in which machine-learning methods are used with available prior data to discover a closure model rather than prescribe one. However, unlike other data-driven approaches, the training data in autonomic closure is obtained internally at a test-filter scale at each point and time in the simulation it-

self, rather than being provided separately from prior simulations or experiments. Importantly, autonomic closure is not a closure model; instead it is a closure methodology that enables essentially model-free “on the fly” closure of any subgrid term.

The need in fully dynamic implementations of autonomic closure to solve a local system identification problem at each point and time in a simulation can make its computational cost far higher than that of traditional prescribed closure models. That is certainly the case when the number of degrees of freedom in the generalized representation of the subgrid terms is large; e.g., the implementation in Ref. [59] involved nearly 6000 degrees of freedom in its generalized representation. Some additional computational cost is acceptable in order to gain the increased accuracy in fields such as $\tau_{ij}(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$ that autonomic closure provides, as seen in Figure 1, since subgrid stress evaluation is typically only a small fraction of the total computational cost of a simulation. However, for the implementation in Ref. [59] the subgrid stress evaluation was $\mathcal{O}(10^4)$ more costly than for traditional prescribed closure models. This cost must be reduced by several orders of magnitude, as has been done in the present study, to make autonomic closure practical for LES.

1.7 Present Study

The cost of autonomic closure can be controlled by varying the number of degrees of freedom in the underlying nonparametric relation and by other choices in its implementation, though these choices can affect the accuracy of the resulting τ_{ij} and $P(\mathbf{x}, t)$ fields. Therefore the main issue addressed in the present study is whether there are implementations of autonomic closure that are efficient enough to be practical for LES while retaining the accuracy in τ_{ij} and $P(\mathbf{x}, t)$ seen in Figure 1 from the computationally costly implementation in Ref. [59]. To address this, results are presented from *a priori* tests that quantify the effects of various implementation choices in autonomic closure. In particular, comparisons are presented of autonomically determined τ_{ij} and $P(\mathbf{x}, t)$ fields with corresponding true subgrid stress and subgrid production fields to find implementations that are both efficient and accurate.

Of key interest is whether large forward and backward scatter in the production fields from

such efficient implementations of autonomic closure remain at the right magnitudes in regions at the right locations and having the right sizes and shapes. To evaluate this, metrics are developed and applied that quantify how the resulting spatial support on which large production values are concentrated compares with corresponding true $P(\mathbf{x}, t)$ fields. From this highly efficient implementations of autonomic closure are identified that remain nearly as accurate as that in Ref. [59] but at computational costs that are $\mathcal{O}(10^5)$ smaller. These implementations are accurate and efficient enough for practical use in large eddy simulations, allowing future *a posteriori* testing of this new closure methodology.

1.7.1 Objectives

The present study seeks to develop accurate and efficient implementations of autonomic closure that are sufficient to allow this new general closure methodology to be practically implemented in large eddy simulations. In particular, it seeks to identify generalized representations for the subgrid stress that are far more computationally efficient than that used in Refs. [58, 59], while retaining essentially the same accuracy that was obtained in these early implementations of autonomic closure.

Specifically, this study first examines efficiencies that are available within the original series-based generalized representation of the subgrid stress tensor in terms of velocities and pressures in the simulation [58, 59]. It identifies effective simplifications both in the series representation itself and in how the series representation can be implemented within the autonomic closure methodology. Reductions of several orders of magnitude in the computational cost of autonomic closure, with essentially no loss of accuracy, will be seen to be possible through these simplifications of the original series-based representation.

The study then goes further by using key concepts from “representation theory” to consider generalized frame-invariant representations for the subgrid stress tensor. The resulting representations will be seen to offer even far greater efficiencies in the autonomic closure methodology than was possible with simplifications of the original non-frame-invariant series representation. Reductions of up to five orders of magnitude in the computational cost of autonomic

closure will be seen to be possible with certain of these frame-invariant tensor representations, with essentially no loss in accuracy.

Two types of frame-invariant generalized representations are investigated here. The first type expresses the subgrid stress directly in terms of frame-invariant outer products of the velocities available in the simulation. This has certain similarities to the original truncated series representations in Refs. [58, 59], but it is a provably complete representation that does not require any truncation, and is far more compact than the original series representations. The second type seeks an even more compact representation in terms of frame-invariant combinations of the strain rate tensor, the rotation rate tensor, and gradients of these tensors. In principle, these representations contain precisely the same “amount” of information as do the frame-invariant representations in terms of outer products of the velocities. However, because they are even more compact they provide correspondingly fewer degrees of freedom across which the autonomic closure methodology can optimize the local stress representation, which can lead to reductions in the accuracy of the results.

In general, throughout this work it is this interplay between efficiency and accuracy that is central to determining the most “accurate and efficient” generalized representation and its implementation within the autonomic closure methodology. The present study identifies such a recommended “best” generalized representation and implementation, and shows that the resulting accuracy in representing the subgrid stress fields and the subgrid production field is far higher than any traditional prescribed closure model. Moreover, this representation and its implementation are shown to be efficient enough to be used in practical large eddy simulations.

1.7.2 Organization of the Dissertation

Analyses, results, and conclusions that address the objectives of this study are organized in this dissertation as follows.

- Chapter 2 first provides a detailed description of the autonomic closure methodology and various choices that can be made in implementing it.

- Chapter 3 then develops a set of suitable metrics that are used throughout this study to assess the accuracy of various implementations of autonomic closure as well as of traditional prescribed closure models. These include metrics for quantifying the scale-dependent structure of support-density fields on which large positive and negative values of $P(\mathbf{x}, t)$ are concentrated.
- Chapter 4 then applies these metrics to quantify the accuracy and efficiency of various implementations of nonparametric series-based generalized representations in autonomic closure, and identifies the series-based implementation that provides the best balance of accuracy and efficiency.
- Chapter 5 then uses representation theory to consider all possible frame-invariant generalized representations for the subgrid stress, including those formed nonparametrically in terms of the resolved-scale velocities and those that are formed parametrically in terms of derived quantities such as the strain rate tensor, the rotation rate tensor, and their gradients.
- Chapter 6 then uses the same metrics to assess the accuracy and efficiency of these frame-invariant generalized representations. It then identifies the “best” representation and its associated implementation, and compares results from this most accurate and efficient implementation of autonomic closure with results from traditional prescribed closure models.
- Chapter 7 then summarizes major conclusions from this study and discusses their implications for achieving computational efficiency and accuracy across essentially all resolved scales in multi-physics large eddy simulations via static or dynamic implementations of autonomic closure.

THE AUTONOMIC CLOSURE METHODOLOGY

This chapter presents a complete description of the autonomic closure methodology [58, 59] and the implementation choices within it that affect its accuracy and computational cost.

2.1 Traditional Prescribed Closure Models vs. Autonomic Closure

Traditional prescribed closure models represent a subgrid term in a predefined way in terms of specified parameters, such as the resolved strain rate \tilde{S}_{ij} , based on theoretical or other considerations, often with one or more model constants allowed to vary locally via a dynamical procedure [27, 46, 47]. In contrast, autonomic closure as proposed in Refs. [58, 59] is based on a highly generalized nonparametric representation of subgrid terms using only the primitive variables in a simulation. This generalized nonparametric representation may have $\mathcal{O}(10^2 - 10^4)$ degrees of freedom that are determined dynamically at every point \mathbf{x} and time t in the simulation. Such a highly generalized representation in the primitive-variable values on a local set \mathbf{S} of stencil points around each space-time point (\mathbf{x}, t) removes the need for a predefined parametric model. Doing so allows far greater adaptability of the subgrid closure to the local turbulence state than is possible even with dynamical forms of traditional prescribed closure models.

2.2 Generalized Representations at the Test and LES Scales

Here the autonomic closure methodology is applied to the subgrid stress $\tau_{ij}(\mathbf{x}, t)$. Although the methodology can be generalized to multi-time stencils [58, 59], the present study considers a time-local implementation. An underlying general nonparametric representation F_{ij} for the local subgrid stress τ_{ij} can then be expressed in the primitive variables $\tilde{\mathbf{u}}$ and \tilde{p} as

$$\tau_{ij}(\mathbf{x}) \approx \tau_{ij}^F(\mathbf{x}) \equiv F_{ij}[\tilde{\mathbf{u}}(x + x'), \tilde{p}(\mathbf{x} + \mathbf{x}') \forall \mathbf{x}' \in \tilde{\mathbf{S}}], \quad (2.1)$$

where $\tilde{\mathbf{S}}$ is a set of points \mathbf{x}' that define a stencil (here $3 \times 3 \times 3$) with separation $\tilde{\Delta}$ on the LES grid, as shown in Figure 2. Equation (2.1) is a general nonparametric relation between the local subgrid stresses $\tau_{ij} \equiv \widetilde{u_i u_j} - \tilde{u}_i \tilde{u}_j$ and the resolved-scale variables $\tilde{\mathbf{u}}$ and \tilde{p} on the LES grid. F_{ij} in (2.1) should ideally reflect the local turbulence state at \mathbf{x} , including nonlinear, nonlocal, nonequilibrium and other effects that determine how the local $\tau_{ij}(\mathbf{x})$ can be best obtained from the resolved-scale values $\tilde{\mathbf{u}}$ and \tilde{p} at the stencil points $\tilde{\mathbf{S}}$ centered on \mathbf{x} .

Analogous to the subgrid stresses τ_{ij} , we consider local test stresses $T_{ij} \equiv \widehat{u_i u_j} - \widehat{u}_i \widehat{u}_j$ [27–29, 39, 40, 46–54], which can be obtained from the resolved velocities $\tilde{\mathbf{u}}$ by applying a test filter $(\widehat{\cdot})$ having a larger length scale $\widehat{\Delta} = \alpha \tilde{\Delta}$. If F_{ij} reflects the local turbulence state at \mathbf{x} then, in the same way that τ_{ij} is related to the resolved-scale variables $\tilde{\mathbf{u}}$ and \tilde{p} on the set $\tilde{\mathbf{S}}$ of LES-scale stencil points centered on \mathbf{x} with separation $\tilde{\Delta}$, so also should T_{ij} be related to the test-scale variables $\widehat{\mathbf{u}}$ and \widehat{p} on the corresponding set $\widehat{\mathbf{S}}$ of test-scale stencil points centered on \mathbf{x} with separation $\widehat{\Delta}$. In other words, analogous to (2.1), this requires

$$T_{ij}(\mathbf{x}) \approx T_{ij}^F(\mathbf{x}) \equiv F_{ij}[\widehat{\mathbf{u}}(\mathbf{x} + \mathbf{x}'), \widehat{p}(\mathbf{x} + \mathbf{x}') \forall \mathbf{x}' \in \widehat{\mathbf{S}}], \quad (2.2)$$

where the stencil $\widehat{\mathbf{S}}$ is the same as stencil $\tilde{\mathbf{S}}$ but is defined on the test-scale grid, as shown in Figure 2.

The central idea in autonomic closure is that at each point \mathbf{x} and time t , the known test stress value $T_{ij}(\mathbf{x})$ and the known surrounding test-scale variables $\widehat{\mathbf{u}}$ and \widehat{p} on the stencil $\widehat{\mathbf{S}}$ in (2.2) can be used to obtain information about the local form of $F_{ij}(\mathbf{x})$. Repeating this at multiple training points within a bounding box centered on \mathbf{x} allows $F_{ij}(\mathbf{x})$ to be determined sufficiently accurately that it can be generalized, in a machine learning sense, from the test scale to the LES scale. The resulting $F_{ij}(\mathbf{x})$ is then used to determine $\tau_{ij}(\mathbf{x})$ via (2.1) from the surrounding variables $\tilde{\mathbf{u}}$ and \tilde{p} on the stencil $\tilde{\mathbf{S}}$ centered on \mathbf{x} . The relation in (2.1) and (2.2) has a sufficient number N of degrees of freedom that it is free to adapt to the local turbulence state at \mathbf{x} to make $T_{ij}^F \approx T_{ij}$, and thereby make $\tau_{ij}^F \approx \tau_{ij}$.

It is in this way that autonomic closure accesses internal training data within the simulation, namely the test stresses within a bounding box centered on \mathbf{x} , to discover the local connection

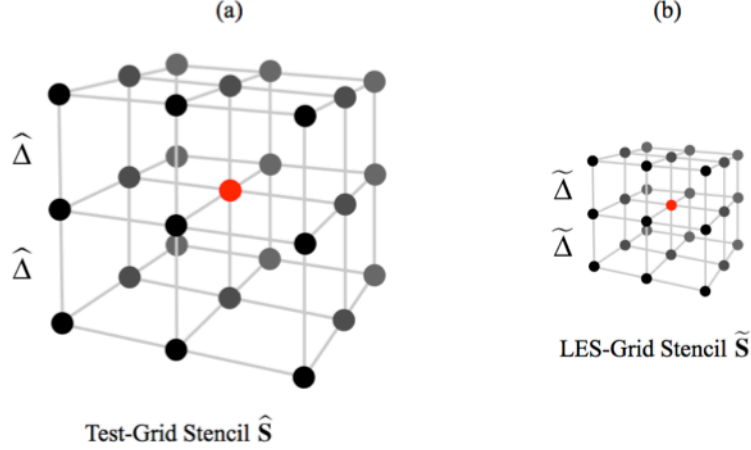


Figure 2. Stencils $\widehat{\mathbf{S}}$ and $\widetilde{\mathbf{S}}$ each centered on point-of-interest \mathbf{x} (red dot); (a) test-grid stencil $\widehat{\mathbf{S}}$ on which \mathbf{h}_{ij} is determined, (b) LES-grid stencil $\widetilde{\mathbf{S}}$ on which resulting \mathbf{h}_{ij} is used to evaluate $\tau_{ij}^{\mathcal{F}}(\mathbf{x})$

$F_{ij}(\mathbf{x})$ between the local test stress $T_{ij}(\mathbf{x})$ and the local test-scale primitive variables $\widehat{\mathbf{u}}$ and \widehat{p} . It then uses this $F_{ij}(\mathbf{x})$ to evaluate the local subgrid stress $\tau_{ij}(\mathbf{x})$ from the local resolved-scale primitive variables $\widetilde{\mathbf{u}}$ and \widetilde{p} . In effect the simulation itself provides the training data, which are used to solve the local, nonlinear, nonparametric system identification problem that discovers the local connection between the subgrid stress and the resolved-scale primitive variables.

Galilean invariance is enforced by subtracting the velocity at the stencil center point from the velocities on the stencil. F_{ij} in (2.1) and (2.2) could be represented in a form that explicitly imposes the tensor invariance and realizability properties of τ_{ij} , as is done in Chapter 5. For example, F_{ij} could be represented in the tensor integrity basis [67, 68, 72–74] for the strain rate and rotation rate tensors, $\widetilde{\mathbf{S}}_{ij}$ and $\widetilde{\mathbf{R}}_{ij}$, which is the basis for many prescribed subgrid stress models. However, that presumes that τ_{ij} depends only on combinations of these two tensors, whereas the nonparametric formulation in (2.1) in the primitive variables $\widetilde{\mathbf{u}}$ and \widetilde{p} makes no such assumption. Instead, as is common in many machine learning methods [67, 68], the tensor invariance and realizability properties are inherent in the training data T_{ij} , which informs the learned F_{ij} in (2.2) and thereby implicitly communicates these properties to τ_{ij} in (2.1). Results will be shown in Chapter 4 demonstrating that $\tau_{ij}(\mathbf{x}, t)$ and $P_{ij}(\mathbf{x}, t)$ from a non-frame-

invariance preserving generalized representation F_{ij} in the autonomic closure methodology are far more accurate than those from common traditional prescribed closures that explicitly enforce tensor frame-invariance properties.

2.3 Formulation of the Optimization Problem

Although any sufficiently general form for F_{ij} could be used in (2.1) and (2.2), the first part of this study chooses a Volterra-like series [75] in \mathbf{u} and p , as in Refs. [58, 59], namely the sum of all products of all orders of all variables at all points on the stencil, including all possible multi-point multi-variable products at each order. Such a representation is highly general. Even if truncated after second order with a $3 \times 3 \times 3$ stencil, F_{ij} for each ij consists of $N = 5995$ zeroth-, first-, and second-order products of \mathbf{u} and p in (2.1) and (2.2), each having a separate coefficient $h_{ij}^{(k)}$ with $k = 1, \dots, N$, where N is the number of degrees of freedom in F_{ij} . This set of coefficients for each ij is thus an N -length column vector denoted \mathbf{h}_{ij} . To gain computational efficiency, F_{ij} could be truncated at lower orders, or restricted to single-point products, or limited only to the velocities \mathbf{u} on the stencil. Even if F_{ij} is truncated after first order and limited to single-point velocities on a $3 \times 3 \times 3$ stencil, it still contains $N = 82$ first-order terms in the \mathbf{u} components for each ij , and thus has far more degrees of freedom than do dynamic versions of traditional prescribed closure models.

With such a series for F_{ij} , the stress value τ_{ij} at the stencil center point \mathbf{x} can be written from (2.1) as

$$\tau_{ij}(\mathbf{x}) \approx \tau_{ij}^F(\mathbf{x}) = \tilde{\mathbf{V}}\mathbf{h}_{ij}, \quad (2.3)$$

where $\tilde{\mathbf{V}}$ is the N -length array containing the known values of all products of all orders of $\tilde{\mathbf{u}}$ and \tilde{p} in F_{ij} at all points $\mathbf{x} + \mathbf{x}'$ on the stencil $\tilde{\mathbf{S}}$. If \mathbf{h}_{ij} is known then the value of τ_{ij} at the stencil center point \mathbf{x} can be obtained from (2.3). To determine \mathbf{h}_{ij} , (2.2) can be similarly written as $T_{ij} \approx T_{ij}^F = \hat{\mathbf{V}}\mathbf{h}_{ij}$, where $\hat{\mathbf{V}}$ is the N -length array containing the known values of all products of all orders of the test-filtered primitive variables $\hat{\mathbf{u}}$ and \hat{p} in F_{ij} at all points $\mathbf{x} + \mathbf{x}'$ on the stencil $\hat{\mathbf{S}}$, and where the test stress value T_{ij} at the stencil center point \mathbf{x} is known.

Repeating this with the stencil $\widehat{\mathbf{S}}$ centered at each of M training points within a local bounding box around the point \mathbf{x} , in which variations in the turbulence state embodied in F_{ij} are taken to be negligible, then $\widehat{\mathbf{V}}$ becomes an $M \times N$ matrix. With \mathbf{T}_{ij} denoting the corresponding M -length column vector consisting of the known T_{ij} values at the M training points, we then have for each ij

$$T_{ij}(\mathbf{x}) \approx T_{ij}^F(\mathbf{x}) = \widehat{\mathbf{V}}\mathbf{h}_{ij}. \quad (2.4)$$

The size of the bounding box, typically extending along homogeneous directions, determines the maximum number of available training points within it. The chosen number M of training points and the bounding box volume V_B determine the relative training point spacing $(V_B/M)^{1/3}$, which together with M determines how much effectively independent information is being used to characterize the local turbulence state via F_{ij} , or equivalently via \mathbf{h}_{ij} , within the bounding box around the point \mathbf{x} . The choice of bounding box size and the number of training points M are part of any implementation of autonomic closure. Regardless of the M and N values, since the vector \mathbf{T}_{ij} and the matrix $\widehat{\mathbf{V}}$ are known, the system in (2.4) may be solved by any number of means. The present study uses a damped least-squares solution [59] of the form

$$\mathbf{h}_{ij} = \left(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I} \right)^{-1} \widehat{\mathbf{V}}^T \mathbf{T}_{ij}. \quad (2.5)$$

where λ is the damping coefficient. When $M/N \gg 1$ the value of λ is unimportant and is set to $\lambda = 10^{-3}$; when $M/N \leq O(1)$ then λ is set to $\lambda = 10^{-1}$. Once the coefficients \mathbf{h}_{ij} at \mathbf{x} have been determined via (2.5), they are used in (2.3) to evaluate at \mathbf{x} .

Note the resulting local \mathbf{h}_{ij} is used only once to evaluate $\tau_{ij}(\mathbf{x}, t)$ at the bounding box center point \mathbf{x} for the current time t . A new set of coefficients \mathbf{h}_{ij} is obtained for each (\mathbf{x}, t) in the simulation. As a result, autonomic closure does not provide a fixed set of coefficients \mathbf{h}_{ij} and thus does not provide a closure model for τ_{ij} . Instead, it is the autonomic methodology itself that is the closure for $\tau_{ij}(\mathbf{x}, t)$.

2.4 Implementation Choices in Autonomic Closure

Implementing autonomic closure involves choices of F_{ij} , N , M , and V_B that impact the generalizability of \mathbf{h}_{ij} , in a machine learning sense, from the test scale to the LES scale, and also determine the computational cost of this closure methodology. For example, the implementation in Ref. [59] used a second-order, non-colocated, velocity-pressure series for F_{ij} , which provided a large number ($N = 5995$) of coefficients \mathbf{h}_{ij} in F_{ij} , and also used the largest possible bounding box, which allowed a large number ($M = 15,625$) of training points to determine \mathbf{h}_{ij} . Results from that implementation verified that autonomic closure produces subgrid stress fields $\tau_{ij}^F(\mathbf{x}, t)$ and subgrid production fields $P^F(\mathbf{x}, t) \equiv -\tau_{ij}^F \tilde{S}_{ij}$ that represent the true $\tau_{ij}(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$ fields over essentially all resolved scales far more accurately than do existing prescribed closure models, such as the dynamic Smagorinsky model [59]. However while the implementation in Ref. [59] is accurate, as can be seen in Fig. 1, it is too computationally costly in comparison with traditional prescribed closure models to serve as a widely-usable alternative closure for practical applications of LES. Therefore, in the following sections this study evaluates the effects of various implementation choices on the accuracy and computational cost of autonomic closure. Specifically, this study considers specific combinations of the following key implementation choices:

- local vs. nonlocal forms based on the bounding box volume $V_B = (n\hat{\Delta})^3$
- velocity-only vs. velocity-pressure series representations F_{ij}
- colocated vs. non-colocated products on the stencils $\tilde{\mathbf{S}}$ and $\hat{\mathbf{S}}$
- first-order vs. second-order series representations F_{ij}
- varying numbers N of coefficients \mathbf{h}_{ij}
- varying numbers M of training points
- varying training point spacing $(V_B/M)^{1/3}$

These allow implementations ranging from (1) a large, second-order, velocity-pressure, non-colocated, nonlocal formulation with $M/N \gg 1$, which provides a large number N of

degrees of freedom in F_{ij} but is computationally costly, to (2) a minimal first-order, velocity-only, colocated, local formulation with $M/N \ll 1$, which has the lowest computational cost but can be expected to be less accurate. The main objective is to determine which implementations of autonomic closure provide high accuracy in $\tau_{ij}(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$ at acceptable computational cost.

2.5 Anticipated Effects of Implementation Choices

It can be expected that various combinations of these implementation choices will affect both the accuracy and efficiency of the resulting implementation of autonomic closure. For instance, because velocities at distant points affect local pressure values on the stencil, a velocity-pressure implementation adds nonlocality beyond the stencil points in F_{ij} and also increases the number N of degrees of freedom in F_{ij} . Both effects can increase generalizability of the resulting \mathbf{h}_{ij} from the test scale to the LES scale, but at increased computational cost over a velocity-only formulation. Similarly, including non-colocated products of the stencil-point variables or increasing the truncation order of the series for F_{ij} increase N and introduce additional physical information in F_{ij} , but at an increased computational cost that may not be merited. Larger numbers M of training points should also lead to increased accuracy, but for any bounding box volume V_B the spacing $(V_B/M)^{1/3}$ between training points becomes smaller as M increases, thus providing relatively less additional independent training information to determine F_{ij} while increasing the computational cost. These expectations suggest a non-trivial tradeoff between N , M , and V_B .

It can also be expected that large bounding box volumes V_B , which allow larger numbers M of widely-spaced (and thus more independent) training points, will provide greater accuracy. However, large bounding boxes may contain substantially different turbulence states, and thus the training points will lead to coefficients \mathbf{h}_{ij} that pertain, in part, to turbulence states other than that at the point of interest \mathbf{x} . Smaller bounding boxes give a more-local implementation that ensures training points are relevant to the local turbulence state at \mathbf{x} , but inherently limit the number of available training points and their relative independence. This suggests that the

most accurate implementation may involve a tradeoff between increased locality via smaller bounding boxes and increased training information via larger bounding boxes.

In the following chapters, this study develops and applies quantitative assessment metrics to determine implementations of autonomic closure that provide accurate and efficient results for the $\tau_{ij}(\mathbf{x}, t)$ fields, and particularly in the corresponding $P(\mathbf{x}, t)$ fields, yet do so at computational costs that are sufficiently low to enable practical use of autonomic closure in large eddy simulations.

METRICS TO ASSESS ACCURACY OF AUTONOMIC CLOSURE

A priori tests are used in this study to assess the relative accuracy and computational cost of various implementations of autonomic closure. Although *a priori* tests alone cannot demonstrate that a subgrid closure will provide stable computations when implemented in an LES code, such tests are the most direct way of assessing the accuracy of any closure approach in representing subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$ and subgrid production fields $P(\mathbf{x}, t)$. The present assessments are thus essential for understanding the accuracy and computational cost of various implementations of autonomic closure, to enable later *a posteriori* tests of implementations that are found here to be both accurate and efficient.

To achieve such quantitative assessments of the accuracy of various implementations of autonomic closure, this study develops metrics to assess the accuracy of any implementation of autonomic closure. As noted in Section 1.2, in highly intermittent fields such as the subgrid production $P(\mathbf{x}, t)$ it may not be possible to have precise point-by-point agreement between the true field and that resulting from a subgrid stress closure. However a closure should (i) produce $P(\mathbf{x}, t)$ fields that are structurally similar to the corresponding true production fields, with large values of forward and backward scatter concentrated in regions at the same locations and of the same size and shape, (ii) produce similar statistical distributions of positive and negative values as in the true $P(\mathbf{x}, t)$ fields, and (iii) produce $\langle P(\mathbf{x}, t) \rangle \geq \varepsilon$, where ε is the true subgrid production rate. Therefore, in addition to comparing average values and statistical distributions of τ_{ij} and P from autonomic closure to corresponding results from the true subgrid stress and production fields, this study develops metrics to quantitatively compare the scale-dependent support-density fields on which large positive and negative subgrid production values are concentrated in the true $P(\mathbf{x}, t)$ fields and in $P^F(\mathbf{x}, t)$ from autonomic closure. These provide

sensitive measures of how well any closure for τ_{ij} represents the detailed spatial structure of regions in which large forward and backward scatter occur in the subgrid production field.

3.1 Pseudo-LES Fields, Test Fields, and Resulting Stress Fields

Direct numerical simulation (DNS) data from a 1024^3 simulation of homogeneous isotropic turbulence at $Re_\lambda = 433$ from the Johns Hopkins Turbulence Database [76, 77] were used for these *a priori* assessments. Velocity $u_i(\mathbf{x}, t)$ and pressure $p(\mathbf{x}, t)$ fields were first projected on a 256^3 grid having regular grid spacing Δ , on which all results are displayed. A spectrally sharp filter with cutoff at $k_{\tilde{\Delta}} = 40$, well within the inertial range, was applied to produce pseudo-LES fields $\tilde{u}_i(\mathbf{x}, t)$ and $\tilde{p}(\mathbf{x}, t)$ on the 256^3 display grid. This grid accommodates wavenumbers up to $k_{\Delta} = 128$, allowing the nonlinear product fields $\tilde{u}_i\tilde{u}_j(\mathbf{x}, t)$, which have wavenumbers up to $2k_{\tilde{\Delta}} = 80$, to be represented without aliasing. The true subgrid stress fields $\tau_{ij}(\mathbf{x}, t) = \widetilde{u_i u_j} - \tilde{u}_i\tilde{u}_j$ were then constructed on the same 256^3 display grid. A second spectrally sharp test filter with cutoff at $k_{\hat{\Delta}} = 20$, providing test-to-LES filter ratio $\alpha \equiv \hat{\Delta}/\tilde{\Delta} = 2$, produced test fields $\hat{u}_i(\mathbf{x}, t)$ and $\hat{p}(\mathbf{x}, t)$ on the same 256^3 display grid. This grid allows the product fields $\hat{u}_i\hat{u}_j(\mathbf{x}, t)$, which contain wavenumbers up to $2k_{\hat{\Delta}} = 40$, to also be represented without aliasing. The test stress fields $T_{ij}(\mathbf{x}, t) = \widehat{u_i u_j} - \hat{u}_i\hat{u}_j$ were then constructed on the 256^3 display grid.

The test stress fields $T_{ij}(\mathbf{x}, t)$ and the test-filtered velocity $\hat{u}_i(\mathbf{x}, t)$ and pressure $\hat{p}(\mathbf{x}, t)$ fields are the only inputs needed for the autonomic closure methodology in Section 2.1. To determine τ_{ij} at any point \mathbf{x} on the display grid, since $\hat{\Delta}/\Delta = k_{\Delta}/k_{\hat{\Delta}} = 128/20 = 6$ the local test-scale grid consisted of every sixth point along each direction on the 256^3 display grid within the specified bounding box volume centered on \mathbf{x} . Due to periodicity of the underlying DNS data, the largest possible bounding box size n^3 spans the entire test field domain, thus $n^3 = 256^3/6^3 = 44^3$. At the other extreme, the smallest bounding box accommodates just one training point centered on the $3 \times 3 \times 3$ test-grid stencil $\hat{\mathbf{S}}$, thus $n^3 = 3^3$. For each of M equally spaced training points in the bounding box, the values of \hat{u}_i and \hat{p} at each point on the test-grid stencil $\hat{\mathbf{S}}$ centered on that training point provide the inputs for one row of the matrix




Figure 3. Typical subgrid production field $P(\mathbf{x}, t)$, showing (a) true field, (b) result from autonomic closure, (c) true support-density field $G(\mathbf{x}, t)$, and (d) corresponding support-density field from autonomic closure.

$\hat{\mathbf{V}}$ in (2.4), and the test stress value T_{ij} at that training point is the corresponding component of the \mathbf{T}_{ij} vector. For each ij pair the resulting $M \times N$ system in (2.4) is then solved via (2.5) to determine the coefficients \mathbf{h}_{ij} at \mathbf{x} , which then provide τ_{ij}^F via (2.3) from the values of $\tilde{\mathbf{u}}$ and \tilde{p} at each point on the LES-grid stencil $\tilde{\mathbf{S}}$ centered on \mathbf{x} .

3.2 Statistical Comparisons of $\tau_{ij}(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$

For each implementation in Section 3.1 the resulting subgrid stress fields $\tau_{ij}^F(\mathbf{x}, t)$ and production fields $P^F(\mathbf{x}, t) \equiv -\tau_{ij}^F \tilde{S}_{ij}$ from autonomic closure are compared with the corresponding true stress and production fields $\tau_{ij}^F(\mathbf{x}, t)$ and $P(\mathbf{x}, t) \equiv -\tau_{ij} \tilde{S}_{ij}$. These comparisons include probability densities of stresses and production to assess if the implementation produces




Figure 4. Subgrid production fields $P(\mathbf{x}, t)$ (leftmost column) and associated support-density fields $G(\mathbf{x}, t)$ filtered at successive scale ratios $\Delta_\Gamma/\tilde{\Delta}$, showing (top row) true production field and (bottom row) result from autonomic closure.

similar distributions of positive and negative values as in the true $\tau_{ij}^F(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$ fields. This study also compares the average subgrid production $\langle P^F \rangle$ from each implementation to the true value ε .

However, probability densities only give the distributions of magnitudes in these fields, but provide no information about the spatial structure of the true fields and those from autonomic closure. To determine whether large magnitudes of forward and backward scatter in the subgrid production fields $P^F(\mathbf{x}, t)$ are concentrated in regions at the same spatial locations and of the same size and shape as in the true $P(\mathbf{x}, t)$ fields, support-density fields for the subgrid production are obtained as described in Sections 3.3 and 3.4. These are then used to obtain the scale-dependent metrics M_1 and M_2 described in Section 3.5 that quantify how closely spatial structures in which large forward and backward scatter are concentrated in $P^F(\mathbf{x}, t)$ from autonomic closure compare with those in the true subgrid production fields $P(\mathbf{x}, t)$.

3.3 Support Fields for the Subgrid Production $P(\mathbf{x}, t)$

Figure 3a,b show a typical comparison of the subgrid production field $P^F(\mathbf{x}, t)$ from autonomic closure with the corresponding true field $P(\mathbf{x}, t)$. Probability densities of subgrid stress and production are used in Chapter 4 to compare magnitudes in these fields. However Fig. 3 also shows structural similarities in $P(\mathbf{x}, t)$ and $P^F(\mathbf{x}, t)$, even in many of the detailed features

of these fields, including regions where large positive and negative values of P are clustered. Of central importance, large magnitudes in $P^F(\mathbf{x}, t)$ in Fig. 3b are clustered in regions at about the same locations and of about the same size and shape as those in $P(\mathbf{x}, t)$ in Figure 3a.

However, despite the clear similarities in Figure 3a,b in the location, size, and shape of regions where large magnitudes of subgrid production are concentrated, the precise point-by-point *rms* differences between $P(\mathbf{x}, t)$ and $P^F(\mathbf{x}, t)$, scaled by P'_{rms} , are nevertheless of $O(1)$. This is due to the strong intermittency in these fields, which leads to large *rms* differences even if the two fields appear nearly identical at all but the smallest scales. Metrics other than the simple *rms* difference are needed that can quantify the spatial structure of these fields, focusing on the regions in which large production magnitudes are concentrated. Such structural metrics should not primarily address the production values themselves, since these are already compared in the probability densities of P and P^F , but should focus on the structure of the *spatial support* on which large production values are concentrated.

The support of a field is the subset of the domain on which the field values are non-zero. This study defines the support on which large magnitudes of the subgrid production fields $P(\mathbf{x}, t)$ and $P^F(\mathbf{x}, t)$ are concentrated, by thresholding the absolute value of each field at a fixed fraction γ of P'_{rms} . This defines $\Sigma(\mathbf{x}, t)$ as either zero or one, depending on whether the absolute value of the subgrid production is below or above the threshold. Points where $\Sigma = 1$ are on the support of large production magnitudes and those where $\Sigma = 0$ are off the support. Thresholding Σ at $\gamma = 0.75$ provides clear identification of the sensible support on which large magnitudes of the subgrid production fields are concentrated.

3.4 Subgrid Production Support-Density Fields $G(\mathbf{x}, t)$

The support for each of the subgrid production fields can be separated into different scales to allow scale-by-scale comparisons of $P(\mathbf{x}, t)$ and $P^F(\mathbf{x}, t)$. From the support $\Sigma(\mathbf{x}, t)$ we define the corresponding support-density field $G(\mathbf{x}, t)$ as

$$G(\mathbf{x}, t) \equiv \int_V \Sigma(\mathbf{x}', t) \Gamma_{\Delta}(|\mathbf{x} - \mathbf{x}'|) \mathbf{x}', \quad (3.1)$$

where $\Gamma_{\Delta}(|\mathbf{x} - \mathbf{x}'|)$ is a convolution filter kernel with filter length scale Δ_{Γ} . This study uses standard Gaussian filters for Γ_{Δ} in (3.1). Whereas the support $\Sigma(\mathbf{x}, t)$ is a discontinuous binary-valued field, the support-density $G(\mathbf{x}, t)$ is a continuous real-valued field to which standard error measures can be applied. Figures 3c,d show the support-density fields $G(\mathbf{x}, t)$ and $G^F(\mathbf{x}, t)$ corresponding to $P(\mathbf{x}, t)$ and $P^F(\mathbf{x}, t)$ in Figures 3a,b. It is apparent that these G fields accurately identify the locations, sizes, and shapes of the regions in which large subgrid production values are concentrated.

Successive filter length scales Δ_{Γ} in (3.1) allow scale-dependent structure in the support-density fields to be determined. Comparisons at the same filter scale between true production support-density fields and those obtained from the closure allow quantitative assessment of scale-by-scale agreement in these fields. Figure 4 shows an example of such scale-dependent comparisons of the support-density fields $G(\mathbf{x}, t)$ and $G^F(\mathbf{x}, t)$ for the subgrid production fields $P(\mathbf{x}, t)$ and $P^F(\mathbf{x}, t)$ in Figure 3 at successive scale ratios $\Delta_{\Gamma}/\tilde{\Delta}$.

3.5 Support-Density Metrics \mathcal{M}_1 and \mathcal{M}_2

From support-density fields $G(\mathbf{x}, t)$ and $G^F(\mathbf{x}, t)$ for $P(\mathbf{x}, t)$ and $P^F(\mathbf{x}, t)$ as in Figure 4, this study uses two metrics to quantitatively compare their spatial structure. At any scale-ratio $\Delta_{\Gamma}/\tilde{\Delta}$ these metrics are defined as

$$\mathcal{M}_1 \equiv \frac{\langle G'(\mathbf{x}, t)G'^F(\mathbf{x}, t) \rangle_V}{\langle G'(\mathbf{x}, t) \rangle_V^{1/2} \langle G'^F(\mathbf{x}, t) \rangle_V^{1/2}} \quad \text{and} \quad \mathcal{M}_2 \equiv \sqrt{\frac{\langle [G(\mathbf{x}, t) - G^F(\mathbf{x}, t)]^2 \rangle_V}{\langle [G(\mathbf{x}, t)]^2 \rangle_V}} \quad (3.2)$$

where \mathcal{M}_1 is the correlation between the support-density fields $G'(\mathbf{x}, t)$ and $G'^F(\mathbf{x}, t)$, with $G' \equiv G - \langle G \rangle_V$, and \mathcal{M}_2 is the normalized *rms* difference between $G(\mathbf{x}, t)$ and $G^F(\mathbf{x}, t)$. The volume averages are over the entire domain. Note $\mathcal{M}_1 \rightarrow 1$ as the two support-density fields become perfectly correlated, and $\mathcal{M}_2 \rightarrow 0$ as the two support-density fields become identical. The variations in \mathcal{M}_1 and \mathcal{M}_2 with scale-ratio $\Delta_{\Gamma}/\tilde{\Delta}$ allow quantitative comparisons of the spatial support-densities on which the true production field $P(\mathbf{x}, t)$ and its associated $P^F(\mathbf{x}, t)$ from the τ_{ij} closure are concentrated.

3.6 Computational Time Scaling

This study also evaluates the computational time for each implementation of autonomic closure. Only the time associated with the subgrid stress evaluation is considered; all other factors are independent of the choice of closure. In autonomic closure there are seven steps that contribute to the computational time, due to operations needed to:

1. Build the $\widehat{\mathbf{V}}$ matrix in Eq. (2.4) from the N degrees of freedom in \mathcal{F}_{ij} for each of the M training points; since $\widehat{\mathbf{V}}$ is an $M \times N$ matrix, the operations count should scale as $M \cdot N$.
2. Multiply $\widehat{\mathbf{V}}^T \widehat{\mathbf{V}}$ in Eq. (2.5); since $\widehat{\mathbf{V}}$ and $\widehat{\mathbf{V}}^T$ are, respectively, $M \times N$ and $N \times M$ matrices, the operations count should scale as $N^2 \cdot M$.
3. Compute the inverse $(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I})^{-1}$ in Eq. (2.5); since $(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I})$ is an N^2 matrix, the operations count should scale as N^3 .
4. Multiply $\widehat{\mathbf{V}}^T \mathbf{T}_{ij}$ in Eq. (2.5); since $\widehat{\mathbf{V}}^T$ is an $N \times M$ matrix and \mathbf{T}_{ij} is an M -length column vector, the operations count should scale as $M \cdot N$.
5. Multiply $(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I})^{-1}$ by $\widehat{\mathbf{V}}^T \mathbf{T}_{ij}$; since $(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I})^{-1}$ is an N^2 matrix and $\widehat{\mathbf{V}}^T \mathbf{T}_{ij}$ is an N -length column vector, the operations count should scale as N^2 .
6. Build the $\widetilde{\mathbf{V}}$ matrix in Eq. (2.3) from the N degrees of freedom in \mathcal{F}_{ij} ; since $\widetilde{\mathbf{V}}$ is a $1 \times N$ array, the operations count should scale as N .
7. Compute τ_{ij} from \mathbf{h}_{ij} and $\widetilde{\mathbf{V}}$ via Eq. (2.3); since \mathbf{h}_{ij} and $\widetilde{\mathbf{V}}$ are, respectively, $N \times 1$ and $1 \times N$ arrays, the operations count should scale as N .

Since N is common to all operations counts, the total count is driven by steps 2 and 3, regardless how large or small M is relative to N . The total operations count then scales as $N^2 \cdot M + N^3$, so the computational time T should scale as

$$T \sim N^3 \left[1 + (M/N) \right] \quad (3.3)$$

This suggests that there may be little reduction in total computational cost from using $M \ll N$ training points, which, in turn, might suggest using all training points available in the bounding box. However, there is diminishing benefit from increasing M once the training point spacing $(V_B/M)^{1/3}$ has become so small that little additional independent training information is being gained.

Table 1. Implementations of autonomic closure considered in Chapter 4, showing case number and code, number N of degrees-of-freedom in F_{ij} , relative bounding box size n^3 , number M of training points in bounding box, number of training points per degree-of-freedom M/N , relative training point spacing, computational time, and ratio of autonomic and true average subgrid dissipation. Case codes: First character N: non-colocated, C: colocated; Second character L: local, G: nonlocal; Third character 1: first-order, 2: second-order; Primes: velocity-pressure cases; DS = dynamic Smagorinsky model, BD = Bardina scale similarity model; + from (4.1), actual times 10-20X longer due to memory management for large matrix operations.

Case	Code	N	$n^3 = \frac{V_B}{\Delta^3}$	M	M/N	$\frac{(V_B/M)^{1/3}}{\Delta}$	t(s)	$\frac{\langle P^F \rangle}{\langle P \rangle}$
1a	CL14	82	7^3	328	4.0	1.0	45	1.18
1b	CL14'	109	8^3	436	4.0	1.1	95	0.99
2a	CL18	82	9^3	656	8.0	1.0	85	1.15
2b	CL18'	109	10^3	872	8.0	1.0	152	1.13
3a	CL24	244	10^3	976	4.0	1.0	624	1.08
3b	CL24'	379	12^3	1516	4.0	1.0	2530	1.00
4a	CL28	244	13^3	1952	8.0	1.0	1186	1.06
5a	NG2	3403	44^3	17576	5.2	1.7	130,821+	1.19
5b	NG2'	5995	44^3	17576	2.9	1.7	261,543+	1.23
6a	CG24	244	44^3	976	4.0	4.4	624	1.04
6b	CG28	244	44^3	1952	8.0	3.5	1186	0.99
7a	CL1(3)	82	3^3	27	0.2	1.1	15	0.68
7b	CL2(3)	244	3^3	27	0.07	1.1	114	0.68
8a	CL1(5)	82	5^3	64	0.9	1.2	18	0.54
8b	CL2(5)	244	5^3	64	0.3	1.2	130	0.88
DS	—	—	—	—	—	—	70	1.05
BD	—	—	—	—	—	—	33	-1.19

PERFORMANCE OF SERIES-BASED REPRESENTATIONS

This chapter applies the metrics developed in Chapter 3 to assess the accuracy and efficiency of series-based representations for F_{ij} , of the type discussed in Chapter 2, for various combinations of the implementation choices discussed in Sections 2.4 and 2.5. It then uses those results to determine which implementation of such a series-based representation of autonomic closure provides the greatest accuracy and efficiency in representing subgrid terms such as the subgrid stress $\tau_{ij}(\mathbf{x}, t)$ and the associated subgrid production $P(\mathbf{x}, t)$.

Table 1 lists relevant parameters discussed in Section 2.4 for each of the series-based implementations of autonomic closure considered in this chapter. Together these allow a systematic assessment of the effects of various implementation choices on the performance of autonomic closure. Cases are grouped into three categories. The first category (Cases 1a-4b in Table 1) primarily examines effects of series truncation order, inclusion of pressure, and the relative training ratio M/N , while the bounding box size n^3 and the number M of training points vary as needed to obtain these M/N values. The second (Cases 5a-6b) primarily examines effects of colocated and non-colocated implementations, with and without pressure, keeping the bounding box at its largest possible size and the relative training ratio M/N roughly comparable to the first category. The third category (Cases 7a-8b) primarily examines highly local implementations and series truncation order, with relative training ratios M/N necessarily low due to the small bounding box sizes.

Figures 6-12 present results from these tests of autonomic closure in homogeneous isotropic turbulence. Each figure shows typical comparisons of the resulting normal and shear stress components $\tau_{ij}^F(\mathbf{x}, t)$ with the corresponding true fields $\tau_{ij}(\mathbf{x}, t)$, and comparisons of the corresponding subgrid production fields $P^F(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$.

4.1 Computational Times

The resulting single-core computational time for each implementation of autonomic closure is given in Table 1. This study used standard LAPACK routines that take advantage of symmetry to speed the inversion in (2.5). As noted in Section 3.6, the computational cost for autonomic closure is determined primarily by the time needed to multiply $\widehat{\mathbf{V}}^T \widehat{\mathbf{V}}$, which is expected to scale as $N^2 \cdot M$, and the time needed to compute the inverse in (2.5) for the coefficients \mathbf{h}_{ij} , which is expected to scale as N^3 . In comparison, the time for all other steps in Section 3.6 should be negligible. This suggests that the computational times in Table 1 would be expected to scale as

$$T \sim N^3 \left[1 + (M/N) \right]. \quad (4.1)$$

It is apparent in Figure 5 that (4.1) provides reasonable scaling of T over nearly six orders of magnitude, and it can therefore be used to understand how the implementation parameters N and M affect the computational time.

Note it is only when $M/N \gg 1$ that the number M of training points significantly affects the computational time, and when $M/N \ll 1$ the computational time simply scales as N^3 . It will be seen in Section 4.5 that maintaining the high accuracy available via autonomic closure requires $M/N \approx O(4)$, with little benefit gained from increasing M/N beyond this. As a result, from (4.1) the computational time for such implementations of autonomic closure scales roughly as $T \sim M \cdot N^2$.

4.2 Local vs. Nonlocal Representations

As noted in Section 2.4, large bounding boxes allow for larger numbers M of widely spaced (and thus more independent) training points, which provide more information for determining F_{ij} via \mathbf{h}_{ij} . However, increasingly larger bounding boxes contain increasingly different local turbulence states than that which applies at the bounding box center point \mathbf{x} , and thus yield coefficients \mathbf{h}_{ij} that are influenced by states other than that at \mathbf{x} . On the other hand, increasingly smaller bounding boxes provide a more-local implementation that assures training points are




Figure 5. Scaling of computational time T for autonomic closure with number N of degrees of freedom in F_{ij} and number M of training points in $\widehat{\mathbf{V}}$ in (2.4) and (2.5), showing theoretical scaling vs. actual computational time for cases in Table.1

relevant to the local turbulence state at \mathbf{x} , but they inherently limit the number M of available training points and their relative independence.

To understand the performance of local (small n^3) and nonlocal (large n^3) implementations of autonomic closure, Figure 13 shows probability densities of a typical subgrid stress component τ_{ij} and the subgrid production P , and the scale-dependent support-density metrics \mathcal{M}_1 and \mathcal{M}_2 comparing the accuracy in the support-densities of the subgrid production fields. The red curves show the performance of local implementations (Cases 3a vs 4a in Table 1) and the blue curves are from nonlocal implementations (Cases 6a vs 6b). Cases 3a and 6a have the same number M/N of training points per degree of freedom, as do Cases 4a and 6b, and all these cases are second-order velocity-only implementations, so these comparisons are indicative of the effects of local versus nonlocal implementations.

In Figure 13 it is apparent in both \mathcal{M}_1 and \mathcal{M}_2 that the local implementations give more accurate results for the support on which large magnitudes of subgrid production are concentrated than do the nonlocal implementations. Support-density correlations \mathcal{M}_1 for the local implementations are as high as 98% and rms errors \mathcal{M}_2 decrease to less than 2.3% at the largest scale-ratio, while for the nonlocal implementations the correlations \mathcal{M}_1 only reach 94% and

rms errors \mathcal{M}_2 decrease only to 4%. The pdfs in Figure 13 show that, for the same M/N value, the local implementations give better representation of τ_{ij} and P , especially in the tails that correspond to large magnitudes, than do the corresponding nonlocal implementations. Moreover, note in Section 3.1 that the computational cost is the same for Cases 3a and 6a, and for Cases 4a and 6b, since M/N is the same for each pair, as discussed in Section 4.1. We conclude that, all other factors being the same, local implementations are more accurate than nonlocal ones.

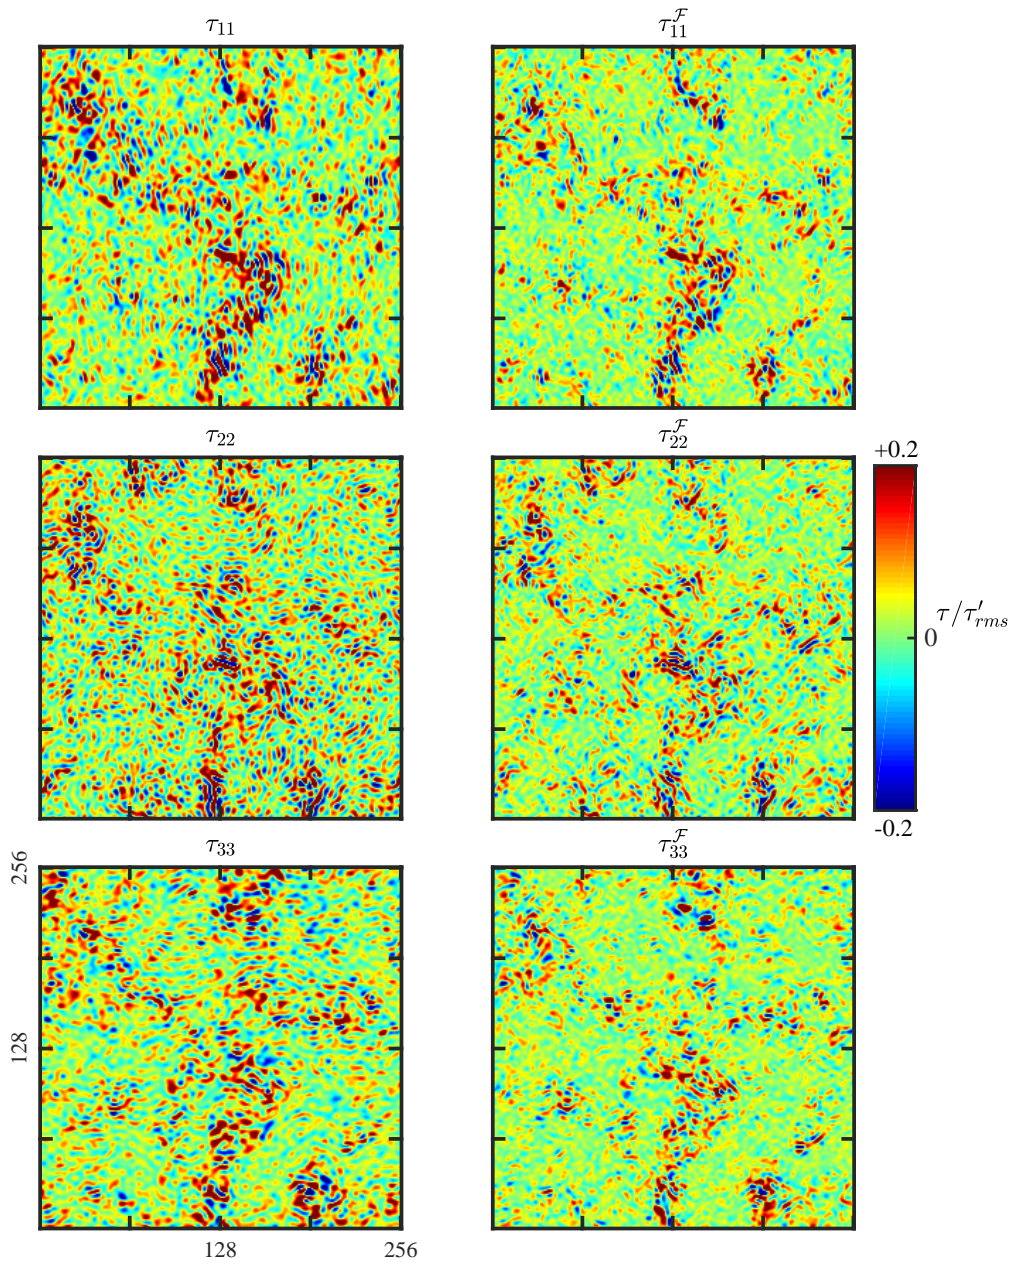


Figure 6 (a). Typical normal subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i = j$




Figure 6 (b). Typical shear subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i \neq j$




Figure 6 (c). Corresponding subgrid production fields $P(\mathbf{x}, t)$.

Figure 6. Local vs nonlocal implementations. Typical comparison of true fields (left) and autonomic closure, case 5a (right), showing (a) normal subgrid stresses, (b) shear subgrid stress, and (c) subgrid production $P(\mathbf{x}, t)$

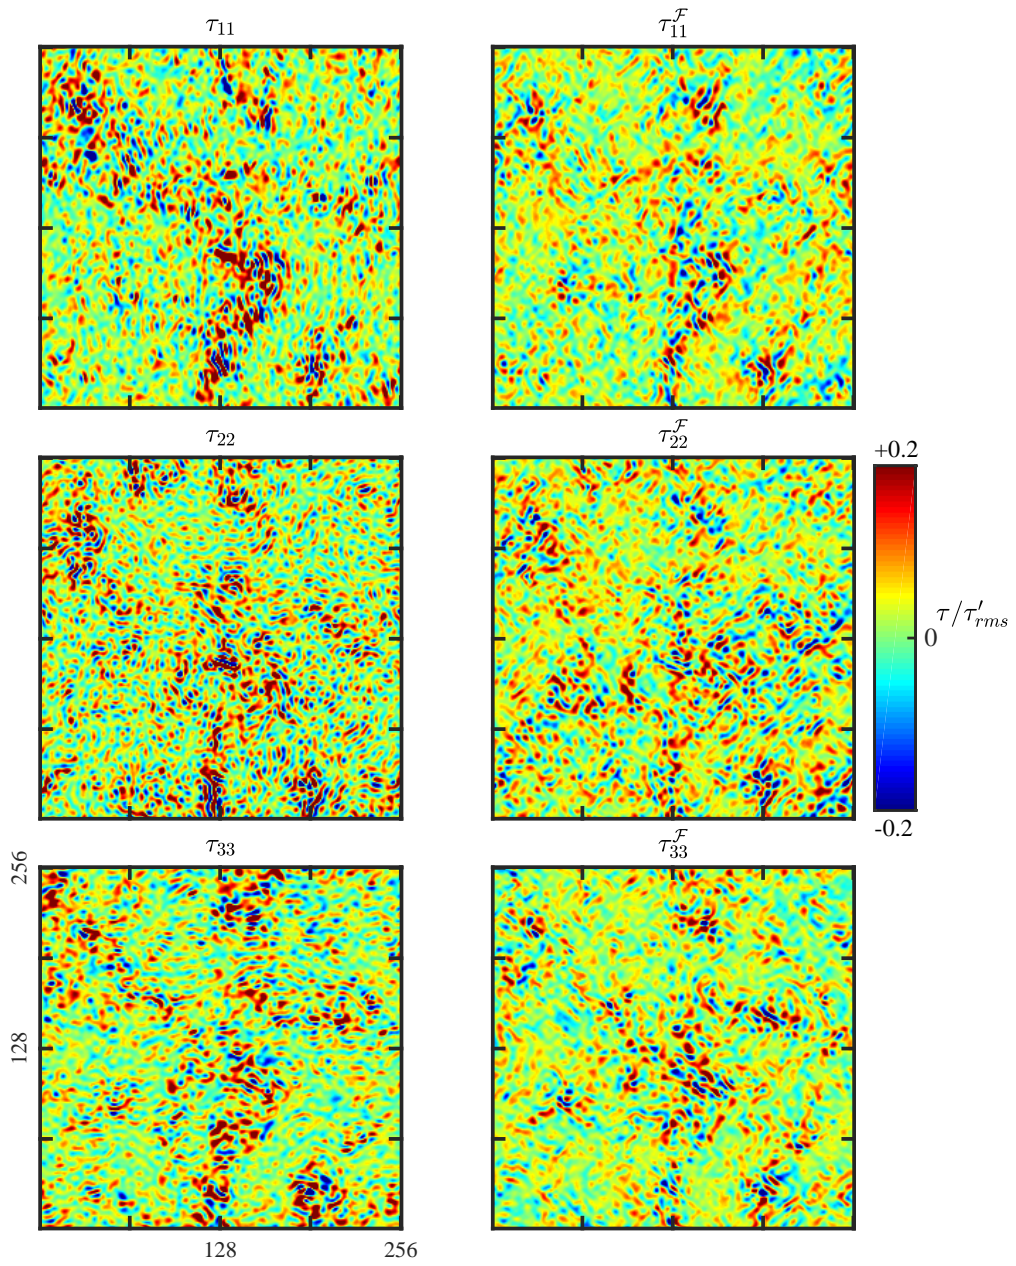


Figure 7 (a). Typical normal subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i = j$

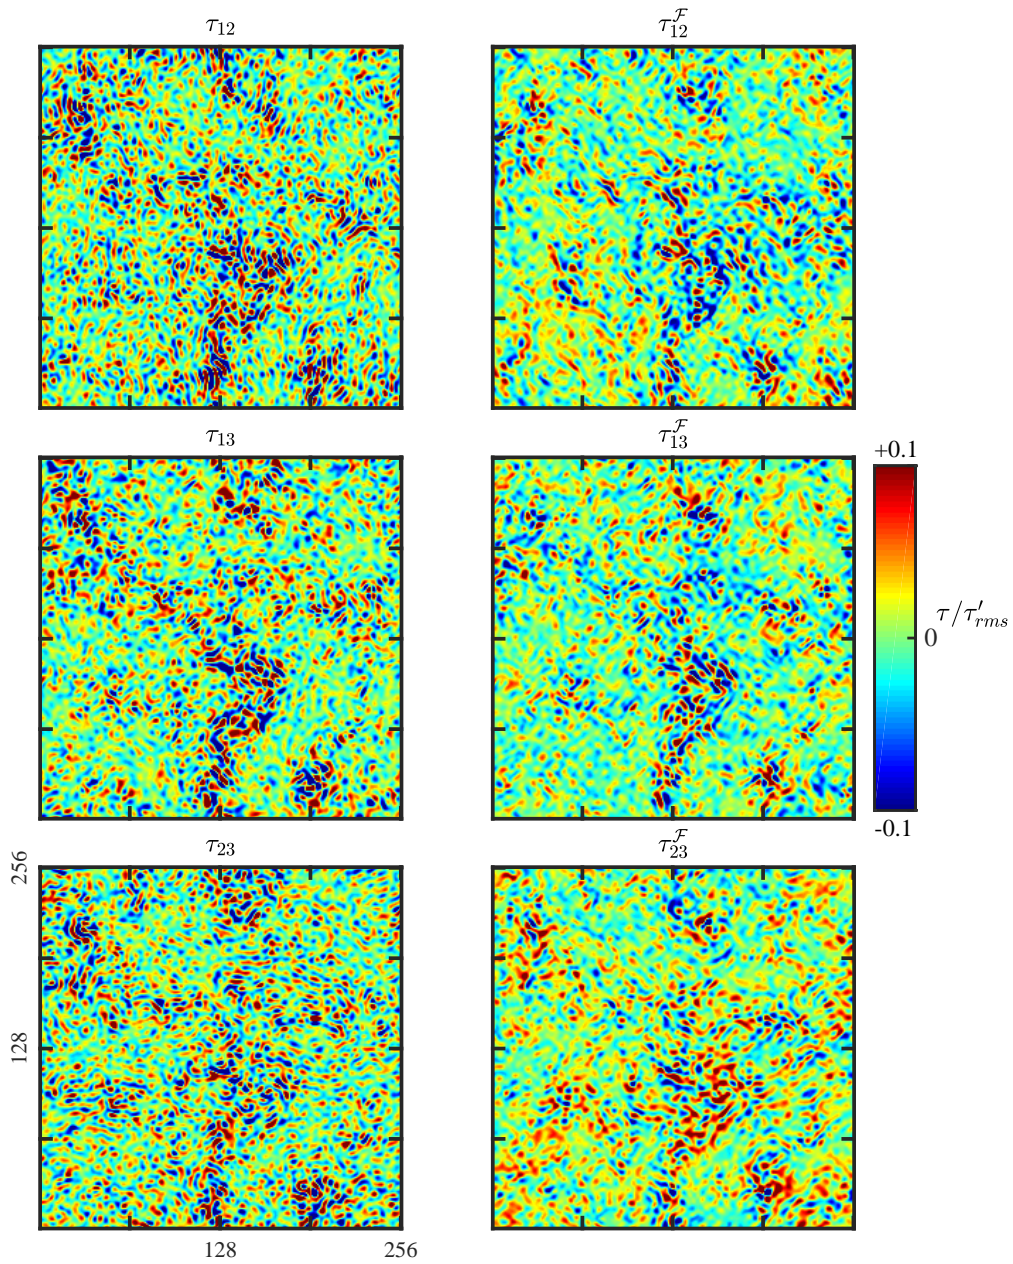


Figure 7 (b). Typical shear subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i \neq j$




Figure 7 (c). Corresponding subgrid production fields $P(\mathbf{x}, t)$.

Figure 7. Local vs Nonlocal implementations. Typical comparison of true fields (left) and autonomic closure, case 6a (right), showing (a) normal subgrid stresses, (b) shear subgrid stress, and (c) subgrid production $P(\mathbf{x}, t)$

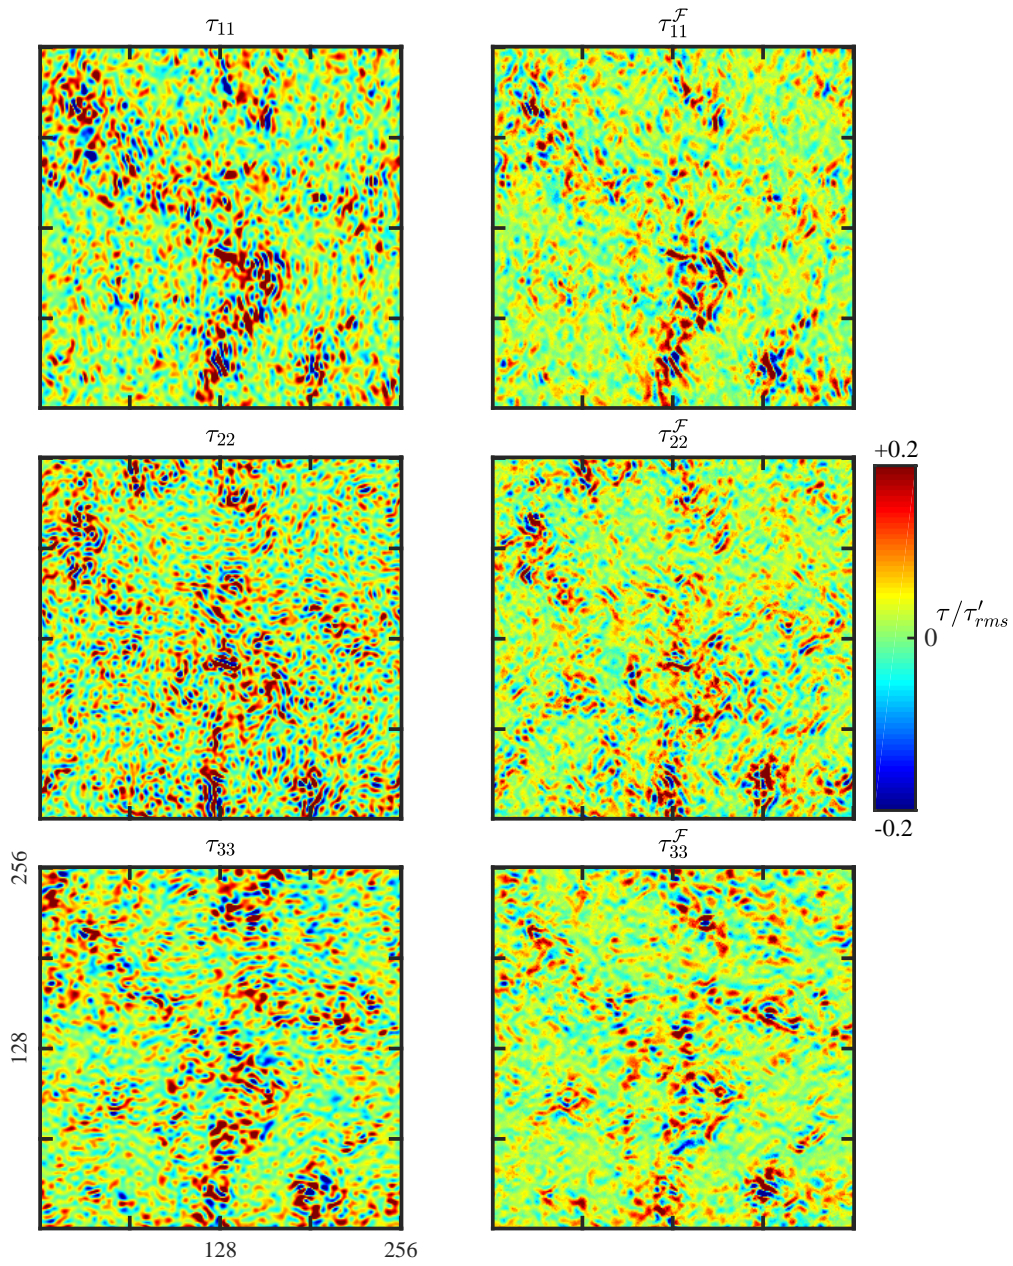


Figure 8 (a). Typical normal subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i = j$




Figure 8 (b). Typical shear subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i \neq j$




Figure 8 (c). Corresponding subgrid production fields $P(\mathbf{x}, t)$.

Figure 8. Velocity-pressure vs Velocity-only implementations. Typical comparison of true fields (left) and autonomic closure, case 4a (right), showing (a) normal subgrid stresses, (b) shear subgrid stress, and (c) subgrid production $P(\mathbf{x}, t)$

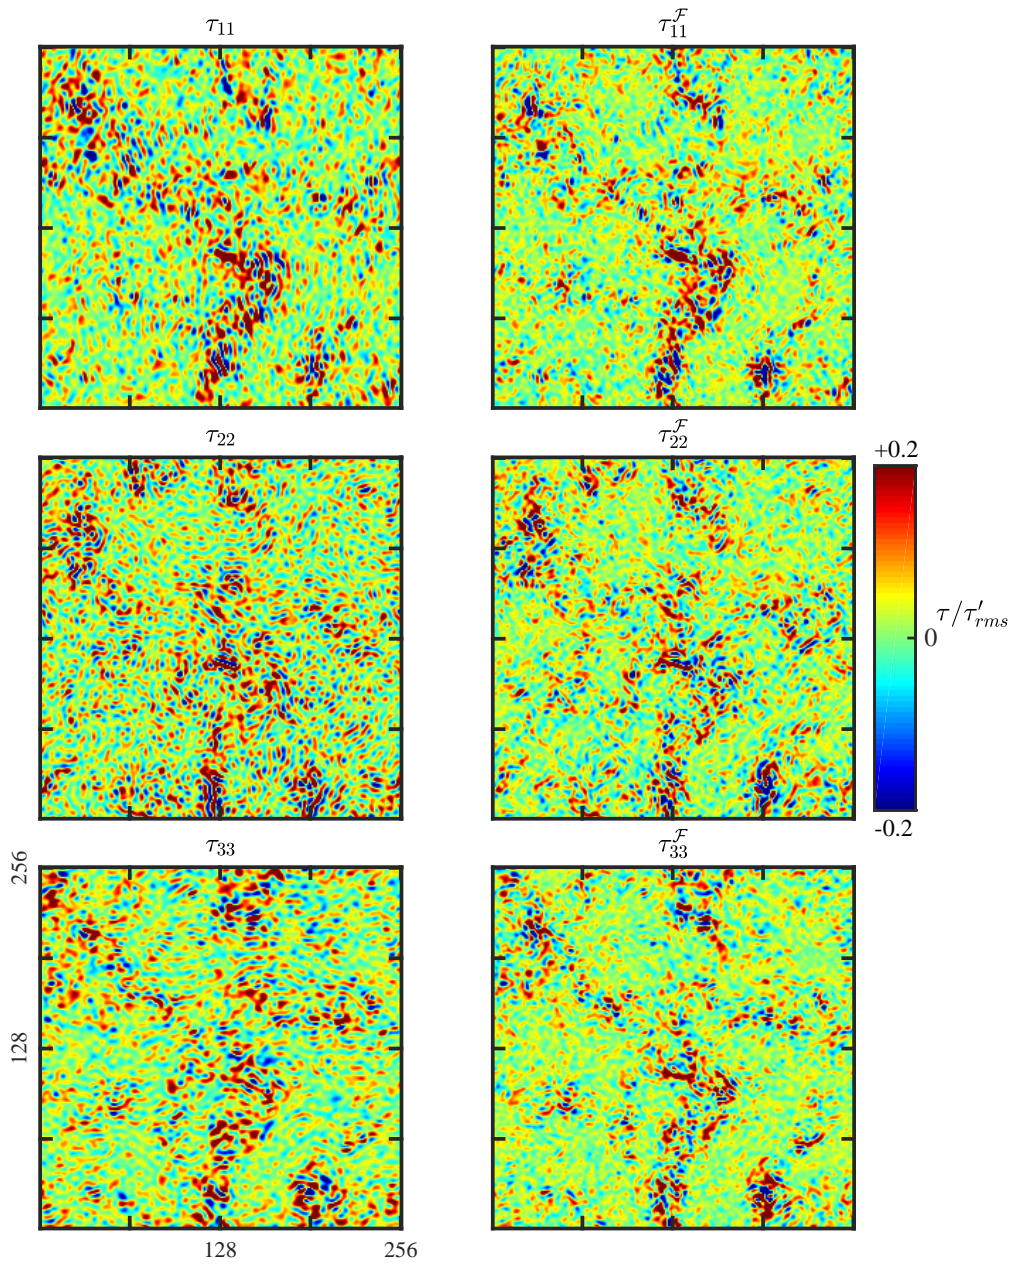


Figure 9 (a). Typical normal subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i = j$




Figure 9 (b). Typical shear subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i \neq j$




Figure 9 (c). Corresponding subgrid production fields $P(\mathbf{x}, t)$.

Figure 9. Velocity-pressure vs Velocity-only implementations. Typical comparison of true fields (left) and autonomic closure, case 5b (right), showing (a) normal subgrid stresses, (b) shear subgrid stress, and (c) subgrid production $P(\mathbf{x}, t)$




Figure 10 (a). Typical normal subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i = j$




Figure 10 (b). Typical shear subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i \neq j$




Figure 10 (c). Corresponding subgrid production fields $P(\mathbf{x}, t)$.

Figure 10. First-order vs second-order implementations. Typical comparison of true fields (left) and autonomic closure, case 1a (right), showing (a) normal subgrid stresses, (b) shear subgrid stress, and (c) subgrid production $P(\mathbf{x}, t)$




Figure 11 (a). Typical normal subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i = j$




Figure 11 (b). Typical shear subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i \neq j$




Figure 11 (c). Corresponding subgrid production fields $P(\mathbf{x}, t)$.

Figure 11. First-order vs second-order implementations. Typical comparison of true fields (left) and autonomic closure, case 1b (right), showing (a) normal subgrid stresses, (b) shear subgrid stress, and (c) subgrid production $P(\mathbf{x}, t)$




Figure 12 (a). Typical normal subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i = j$




Figure 12 (b). Typical shear subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i \neq j$

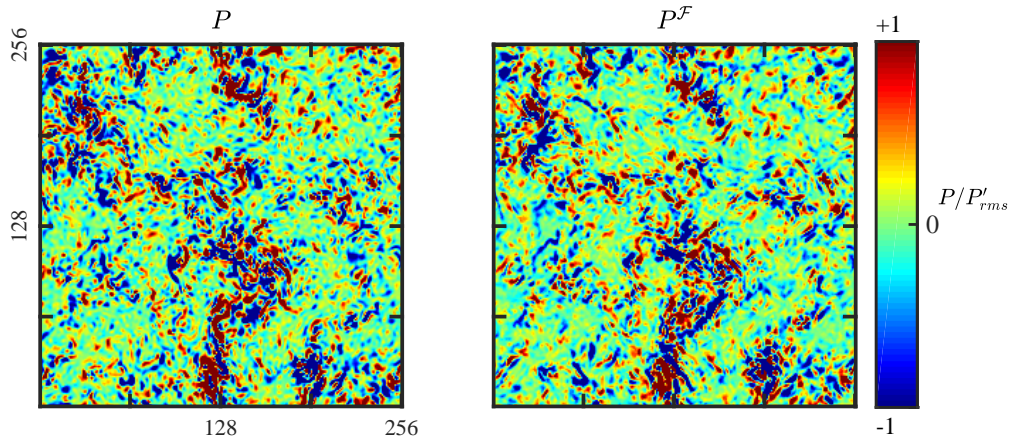


Figure 12 (c). Corresponding subgrid production fields $P(\mathbf{x}, t)$.

Figure 12. Typical comparison of (left column) true subgrid stress component field $\tau_{ij}(\mathbf{x}, t)$ and subgrid production field $P(\mathbf{x}, t)$ with (right column) results from recommended implementation (Case 3a) of autonomic closure; Accuracy is comparable to Figure 1 but at nearly 500X lower computational cost.

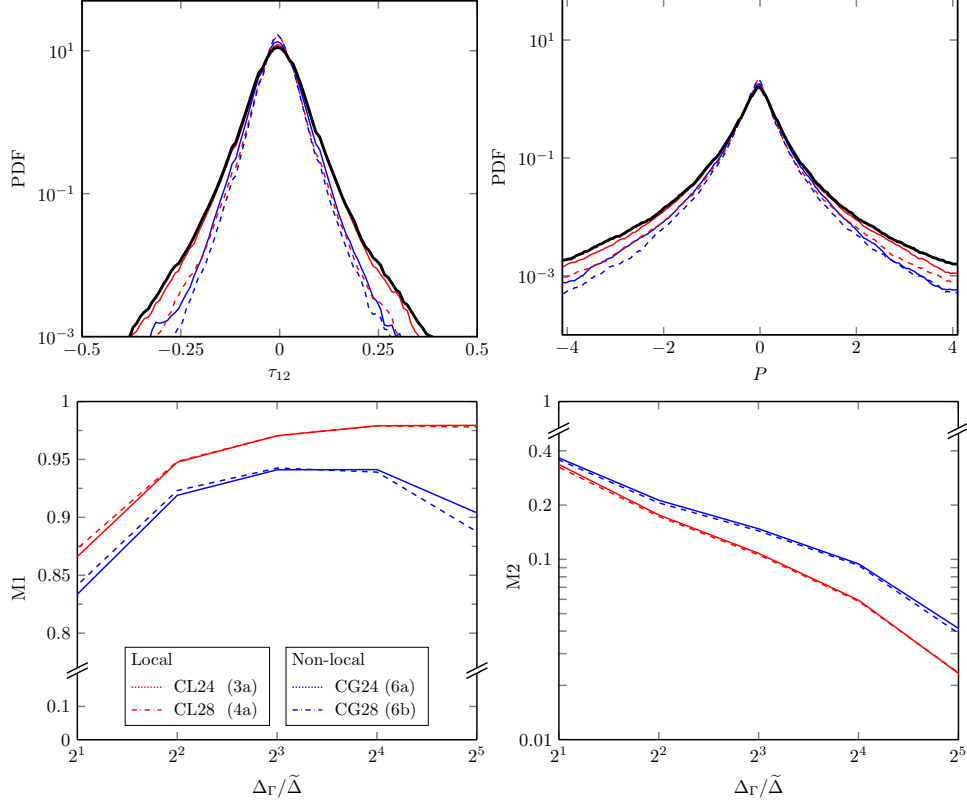


Figure 13. Comparison of nonlocal (blue) and local (red) implementations of autonomic closure; (top) pdfs of (left) typical subgrid stress component τ_{12} and (right) subgrid production P versus (black) exact results, and (bottom) \mathcal{M}_1 and \mathcal{M}_2 variations with scale ratio $\Delta_\Gamma/\tilde{\Delta}$ for subgrid production field.

Figure 13 shows relatively little difference in accuracy between the two local implementations (Cases 3a and 4a) or between the two nonlocal implementations (Cases 6a and 6b). Since each pair differs primarily in its M/N value, there apparently is little benefit in autonomic closure from increasing the number of training points per degree of freedom from $M/N = 4$ to $M/N = 8$. Additionally, the greater separation $(V_B/M)^{1/3}$ between training points in the nonlocal implementations (Cases 6a and 6b) compared to that in the local implementations also appears to provide little benefit. These effects will be examined in greater detail in Section 4.5.

4.3 Velocity-Pressure vs. Velocity-Only Representations

The series used here for F_{ij} in (5) and (6) consists of all possible products up to second order among the variables on a $3 \times 3 \times 3$ stencil, including all multi-variable products. When only

single-point products are included, then F_{ij} has $N = 379$ degrees of freedom for a velocity-pressure implementation and $N = 244$ for a velocity-only implementation. As noted in Section 2.4, including pressure increases nonlocal effects in the nonparametric representation F_{ij} beyond the nonlocality available from the velocities on the stencil. At the same time, a velocity-pressure formulation also increases the number of degrees of freedom N .

Figure 14 compares several velocity-pressure implementations with corresponding velocity-only implementations (Cases 1a vs 1b, 2a vs 2b, 3a vs 3b, 4a vs 4b, and 5a vs 5b). Each pair has the same number M/N of training points per degree of freedom, and in each pair the case number ending in “b” includes pressure. In the \mathcal{M}_1 and \mathcal{M}_2 results, it can be seen that for first-order implementations ($N = 82$ or 109) there is some improvement when pressure is included, but for second-order implementations ($N = 244$ or 379) there is essentially negligible benefit from including pressure. This indicates that the improvement seen in first-order implementations when including pressure is not so much due to the addition of nonlocal effects from the pressure itself, but instead largely due to the greater number N of degrees of freedom in F_{ij} when an additional stencil variable is included.

A velocity-pressure implementation does however lead to greater computational cost over the corresponding velocity-only implementation. As seen in Section 3.1, including pressure typically raises the computational cost by a factor of 2-4. Including pressure may thus be justified for first-order implementations, where the additional degrees of freedom provide some benefit, but not for second-order implementations, where N is already large enough that a further increase in the number of degrees of freedom provides negligible benefits.

Additionally, note in Section 3.1 that noncolocated implementations (Cases 5a and 5b) lead to a large increase in N , and thereby to a large increase in computational time T , but Figure 14 shows that they produce only very little improvement in \mathcal{M}_1 and \mathcal{M}_2 or in the pdfs of τ_{ij} or P . The best colocated implementations perform nearly as well, and do so at far lower computational cost.




Figure 14. Comparison of velocity-pressure (blue) vs. velocity-only (red) implementations of autonomic closure; (top) pdfs of (left) typical subgrid stress component τ_{12} and (right) subgrid production P versus (black) exact results, and (bottom) \mathcal{M}_1 and \mathcal{M}_2 variation in (11a,b) with scale ratio $\Delta_\Gamma/\tilde{\Delta}$.

4.4 First-Order vs. Second-Order Representations

Section 3.1 and (4.1) show there is a substantial increase in computational cost from the increase in N when second-order terms are included in the series for F_{ij} in (2.1) and (2.2). This can be seen by comparing the computational times in Table 1 for Cases 1a vs 3a, Cases 1b vs 3b, Cases 2a vs 4a, and Cases 2b vs 4b. Each pair has the same number M/N of training points per degree-of-freedom. Including second-order terms is seen to increase computational cost by a factor of 14-30 over the cost for a corresponding first-order implementation, consistent with the scaling in (4.1).

Figure 15 compares the performance of first-order and second-order implementations via pdfs of the subgrid stress and production, and the support-density metrics \mathcal{M}_1 and \mathcal{M}_2 for




Figure 15. Comparison of second-order (blue) vs. first-order (red) implementations of autonomic closure; (top) pdfs of (left) typical subgrid stress component τ_{12} and (right) subgrid production P versus (black) exact results, and (bottom) \mathcal{M}_1 and \mathcal{M}_2 variation in (11a,b) with scale ratio $\Delta_\Gamma/\tilde{\Delta}$.

the subgrid production fields. It is evident in the pdfs of τ_{ij} and P and in \mathcal{M}_1 and \mathcal{M}_2 that the second-order (blue) implementations give more accurate results for τ_{ij} and P than do first-order (red) implementations. Part of the benefit from second-order implementations is simply due to the greater number N of degrees of freedom in F_{ij} when second-order terms are retained, consistent with results found in Section 4.3. However, having seen in Section 4.3 that including pressure in a second-order implementation provides negligible benefit, it is apparent that the advantage of these second-order implementations over the corresponding first-order implementations can be understood solely in terms of the velocities on the stencil.

Specifically, the velocities in a first-order implementation restrict F_{ij} to sums and differences of velocity values on the $3 \times 3 \times 3$ stencil. These can account for parametric quantities




Figure 16. Effect of number of training points per degree of freedom in implementations of autonomic closure with $M/N \leq 1$ (blue) and $M/N \gg 1$ (red); (top) pdfs of (left) typical subgrid stress component τ_{12} and (right) subgrid production P versus (black) exact results, and (bottom) \mathcal{M}_1 and \mathcal{M}_2 variation in (11a,b) with scale ratio $\Delta_\Gamma/\tilde{\Delta}$ for subgrid production support-density fields.

such as the strain rate S_{ij} and rotation rate R_{ij} components, as well as their gradients ∇S_{ij} and ∇R_{ij} , all of which alone are only first-order in the velocity components on the stencil. Thus even a first-order velocity-only implementation allows this larger set of parametric quantities to be implicitly represented in F_{ij} than do traditional closures that assume τ_{ij} to depend only on S_{ij} and R_{ij} , with the coefficients \mathbf{h}_{ij} at each point \mathbf{x} determining the relative contributions from each quantity.

A second-order implementation allows an even larger set of such parametric quantities, including tensor products of S_{ij} , R_{ij} , ∇S_{ij} and ∇R_{ij} , to be implicitly represented in F_{ij} . It is this large set of possible tensor products that makes an explicit tensor invariance-preserving parametric formulation of F_{ij} more difficult to implement than is the primitive variable formu-

lation in (2.1) and (2.2), for which tensor invariance properties are instead implicitly inherited in F_{ij} from the test-stress training data, which intrinsically satisfies these properties. Results in Section 4.6 and Section 4.9 show that the present primitive-variable formulation produces far more accurate results for $\tau_{ij}(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$ than do traditional prescribed closure models that explicitly enforce the tensor invariance properties of the subgrid stress.

4.5 Effect of Number of Training Points and Their Separation

As noted in Section 2.4, both the number of training points and their relative independence might be expected to affect the amount of information available in the $\widehat{\mathbf{V}}$ matrix to determine the N degrees of freedom in F_{ij} , and thereby increase its generalizability from the test scale to the LES scale. In general, large numbers M of training points and increased separation $(V_B/M)^{1/3}$ between them could thus lead to improved accuracy in $\tau_{ij}(\mathbf{x}, t)$ and the associated $P(\mathbf{x}, t)$. However, simply increasing both M and $(V_B/M)^{1/3}$ simultaneously by making the bounding box volume V_B large was shown in Section 4.2 to be ineffective, since the resulting less-local implementation then causes $\widehat{\mathbf{V}}$ to contain training points that are not relevant to the turbulence state at the bounding box center point \mathbf{x} . On the other hand, reducing V_B inherently limits the number of available training points and their relative separation.

This study next examines the effect of the number M/N of training points per degree of freedom in F_{ij} . Figure 16 compares the performance of implementations with $M/N \leq 1$ to those having $M/N \gg 1$. In general, cases having $M/N \gg 1$ (red) outperform those with $M/N \leq 1$ (blue). However, it can be seen that there is no apparent benefit from increasing M/N above about 4. For instance, Cases 3a and 4a have nearly identical performance, as do Cases 1a and 2a, even though M/N in each pair are 4.0 and 8.0, respectively. The same can be seen in comparing Cases 6a and 7a in Figure 14. Cases 7a and 8a in Figure 16 also have nearly identical performance, even though their M/N values differ by a factor of four. These results indicate that, as long as M/N is sufficiently larger than one (e.g., $M/N \geq 4$), other aspects of autonomic closure including the locality gained by making the bounding box volume n^3

smaller and allowing for second-order terms in F_{ij} have more effect on the resulting accuracy than does the number of training points per degree of freedom.

Since the number M of training points does not have a primary effect on the resulting accuracy, it may be expected that the spacing between training points, which characterizes their relative independence, also will not have a primary effect. This is supported by comparing the relative training point spacing values $(V_B/M)^{1/3}$ in Section 3.1. For most of these cases, the training point spacing relative to the test-scale grid spacing $\widehat{\Delta}$ does not vary widely, ranging from 1.0-1.2, yet the performance of these implementations varies widely. Cases 6a and 6b, which have spacings roughly 3-4 times larger, show performance no better than many other cases having far smaller training point separation. It is concluded that training point separation by itself does not have a controlling effect on the performance of autonomic closure, as long as the bounding box volume V_B is small enough to provide a local implementation, as noted in Section 4.2.

4.6 Average Subgrid Dissipation Rates

The pdfs of subgrid production in Figs.13b, 14b, 15b, and 16b show that the $P^F(\mathbf{x}, t)$ fields contain large positive and negative values, which combine to produce the average subgrid production rate $\langle P^F \rangle$. For each implementation of autonomic closure in Table 1, the resulting subgrid production fields $P^F(\mathbf{x}, t)$ were averaged to obtain $\langle P^F \rangle$, and the corresponding true subgrid production fields $P(\mathbf{x}, t)$ were averaged to obtain $\langle P \rangle \equiv \varepsilon$. The resulting ratio of autonomic-to-true average subgrid production $\langle P^F \rangle / \varepsilon$ for each implementation is shown in Table 1. Since $\varepsilon > 0$, when $\langle P^F \rangle / \varepsilon > 0$ then $\langle P^F \rangle > 0$, which corresponds to a net average rate of energy transfer from the resolved scales into the subgrid scales. This can be seen in Table 1 to be the case for all these implementations of autonomic closure.

For cases in Table 1 that produce $\langle P^F \rangle / \varepsilon \geq 1$, the closure implementation on average transfers energy out of the resolved scales at a rate equal to or higher than the true subgrid dissipation rate ε . As noted in Section 1.2 this is necessary (but not sufficient) for any closure to maintain computational stability. Most implementations of autonomic closure can be seen in

Table 1 to meet this requirement, with the most obvious exceptions being Cases 7a,b and 8a,b, which all have very small bounding boxes that cause $M/N \leq 1$. Those cases still produce net transfer of energy out of the resolved scales, but at a rate that is far lower than the true subgrid dissipation rate ε . Implementations that produce $\langle P^F \rangle / \varepsilon$ slightly larger than one are desired since they produce net average transfer of energy out of the resolved scales at a rate just slightly higher than the true subgrid dissipation rate ε . As a result such implementations meet the minimum requirement for computational stability while avoiding an excessively large average subgrid production rate that could lead to reduced fidelity in a simulation.

4.7 Results from Highly-Sheared Homogeneous Turbulence

All results presented up to this point in the dissertation regarding the accuracy of subgrid stress fields and subgrid production fields obtained from autonomic closure have been from *a priori* tests based on DNS data for homogeneous isotropic turbulence (HIT). Importantly, there are no assumptions in the autonomic closure methodology that depend in any way on the local state of the turbulence or the particular type of turbulent flow in which this new closure methodology is applied. The only requirement is that the relation between the test stresses T_{ij} and the test-filtered variables in (2.1) must be the same as that between the subgrid stresses τ_{ij} and the LES-filtered variables in (2.2). This should be valid when both the test-filter scale and the LES-filter scale are in the scale-similar inertial range, regardless of the local turbulence state or the particular turbulent flow.

To verify this, in this section results from autonomic closure are presented for subgrid stress fields and subgrid production fields obtained via *a priori* tests based on DNS data for homogeneous sheared turbulence (HST). These tests are based on a highly sheared turbulence case, corresponding to a dimensionless mean shear rate $S^* \equiv (Sk/\varepsilon) = 8.8$. This is typical of the peak mean shear values in the log layer of turbulent boundary layers at moderately high Reynolds numbers. It produces a highly anisotropic turbulent flow, with the anisotropy magnitude and the resulting two-dimensional anisotropy state shown in the turbulent state triangle in Fig. 17.




Figure 17. Two-dimensional anisotropy state, $\text{IIIa} = 0.2686$ and $-\text{IIa} = 0.7876$, (marked in black) characterize highly anisotropic flows ($S^* = 8.8$), typically found in the log-layer ($y^+ \sim 100$) of turbulent boundary layers at moderately high Re .

The DNS data were generated with a code developed by Prof. J. Schumacher at TU-Ilmenau specifically for simulations of homogeneous sheared turbulence. This code uses the usual periodic boundary conditions on the front/back and left/right sides of the domain, but uses slip boundary conditions on the bottom/top surfaces to overcome the large-scale periodic “bursting” cycle that otherwise can make such simulations become unstable. The effects of the slip conditions on the bottom/top surfaces are largely confined near these surfaces, allowing highly sheared homogeneous turbulence to be produced near the center of the domain. The present *a priori* test results were obtained from DNS data planes at the center of the domain.

These DNS data planes were processed for *a priori* testing in exactly the same way as was done for tests with homogeneous isotropic turbulence elsewhere in this dissertation. Indeed, autonomic closure not only requires no knowledge about the mean shear rate, the local turbulence state, or the type of turbulent flow in which it is applied, it does not even have any way of incorporating such information.

The same F_{ij} representation is used (Case 3a) that was found in the previous section to give

accurate and efficient results in *a priori* tests with homogeneous isotropic turbulence. Figures 18-20 present results from these tests of autonomic closure in highly-sheared homogeneous turbulence. Typical comparisons of the resulting normal and shear stress components $\tau_{ij}^F(\mathbf{x}, t)$ with the corresponding true fields $\tau_{ij}(\mathbf{x}, t)$ are shown in Figure 18a and Figure 18b, respectively, and comparisons for the corresponding subgrid production fields $P^F(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$ are shown in Figure 18c. The agreement between the true fields and those obtained from autonomic closure is seen to be similar to what was seen in previous comparisons for homogeneous isotropic turbulence.

Corresponding comparisons of the probability density of subgrid production are shown in Figure 19, where again the agreement can be seen to be similar to what was seen in the comparisons for homogeneous isotropic turbulence.

The subgrid production support-density metrics \mathcal{M}_1 and \mathcal{M}_2 are shown in Figure 20, where the blue curves from these tests in highly sheared homogeneous turbulence are compared with the earlier results from tests in homogeneous isotropic turbulence. It is apparent that autonomic closure represents the regions of large positive and negative subgrid production with similar accuracy for this highly sheared turbulence test case as it did for the isotropic turbulence test case.




Figure 18 (a). Typical normal subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i = j$




Figure 18 (b). Typical shear subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$, $i \neq j$




Figure 18 (c). Corresponding subgrid production fields $P(\mathbf{x}, t)$.

Figure 18. Results for highly sheared homogeneous turbulence, showing typical comparisons of true fields (left) and autonomic closure (Case 3a) (right) for (a) normal stresses, (b) shear stress, and (c) subgrid production $P(\mathbf{x}, t)$




Figure 19. Results for highly-sheared turbulence, showing comparisons of pdfs from true production fields (black) with production fields from autonomic closure (red) and from the dynamic Smagorinsky model (blue).




Figure 20. \mathcal{M}_1 (left) and \mathcal{M}_2 (right) metrics for the production fields from the HIT and the HST cases are near-identical, indicating that the autonomic closure retains structural similarity broadly across the entire spectrum of turbulent flows.

4.8 Recommended Series-Based Representation

Based on the results in Sections 4.1 to 4.7, the most accurate and efficient implementation of autonomic closure among the cases in Section 3.1 is Case 3a. This is a relatively local, second-order, velocity-only, colocated implementation that has been found here to be nearly as accurate as that in Ref. [59] (Case 5b), but at a computational cost that is $\mathcal{O}(10^3)$ smaller.

Typical subgrid stress fields $\tau_{ij}^F(\mathbf{x}, t)$ and subgrid production fields $P^F(\mathbf{x}, t)$ from Case 3a are compared in Figure 12 to corresponding true fields $\tau_{ij}(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$. Of particular importance from the perspective of resolved-scale energetics and computational stability, it is apparent by comparing Figure 12c and Figure 12d that this implementation preserves the structural similarities in $P(\mathbf{x}, t)$ and $P^F(\mathbf{x}, t)$ nearly as well as did the far more computationally costly implementation in Figure 1 (Case 5b). Large positive and negative values in $P^F(\mathbf{x}, t)$ in Figure 12d are clustered in regions at essentially the same locations and of the same size and shape as in $P(\mathbf{x}, t)$ in Figure 12c.

Figure 21 compares PDFs of the subgrid stress fields $\tau_{ij}^F(\mathbf{x}, t)$ from Case 3a with corresponding results for the true fields $\tau_{ij}(\mathbf{x}, t)$. It is apparent that this implementation of autonomic closure accurately produces nearly all aspects of these probability distributions, includ-

ing the tails of these distributions. Figure 22 similarly compares the probability density for the subgrid production fields $P^F(\mathbf{x}, t)$ from Case 3a with the corresponding distribution for the true fields $P(\mathbf{x}, t)$. It is again apparent that this implementation of autonomic closure accurately produces the probability distribution of subgrid production values, including the tails of the distribution that correspond to large positive values (forward) and negative values (backward) scatter in the subgrid production fields. Also, Figures 13 to 16 show that Case 3a provides the most accurate representation of spatial structure in the subgrid production fields, as quantified by \mathcal{M}_1 and \mathcal{M}_2 , among computationally efficient implementations in Section 3.1.

Moreover, Table 1 shows the average subgrid production from Case 3a to be $\langle P^F \rangle \approx 1.08\varepsilon$, indicating that this implementation transfers energy out of the resolved scales at a net average rate just slightly higher than the true average subgrid dissipation rate ε . As a result, this implementation satisfies the minimum requirement for computational stability while avoiding an excessively large average subgrid production rate that could negatively impact fidelity when implemented in a large eddy simulation.

Based on the average subgrid production rate in Table 1 and on the statistical and structural comparisons of $\tau_{ij}^F(\mathbf{x}, t)$ and $P^F(\mathbf{x}, t)$ with $\tau_{ij}(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$ in Figures 13 to 22, the recommended implementation of autonomic closure (Case 3a) satisfies the criteria noted in Section 1.2 for any subgrid closure to be accurate across essentially all resolved scales and to potentially provide computational stability in a simulation. This implementation of autonomic closure is accurate and efficient enough for practical use in large eddy simulations, allowing future *a posteriori* tests to assess its stability when used in an LES code.

4.9 Comparison with Traditional Closure Models

From the preceding sections, the most accurate and efficient implementations of autonomic closure are those based on relatively local, second-order, velocity-only, colocated formulations. Among these, Case 3a was seen in Figures 13 to 22 and Table 1 to give the best results for subgrid stress fields $\tau_{ij}(\mathbf{x}, t)$ and subgrid production fields $P(\mathbf{x}, t) \equiv -\tau_{ij}\tilde{S}_{ij}$ across essen-

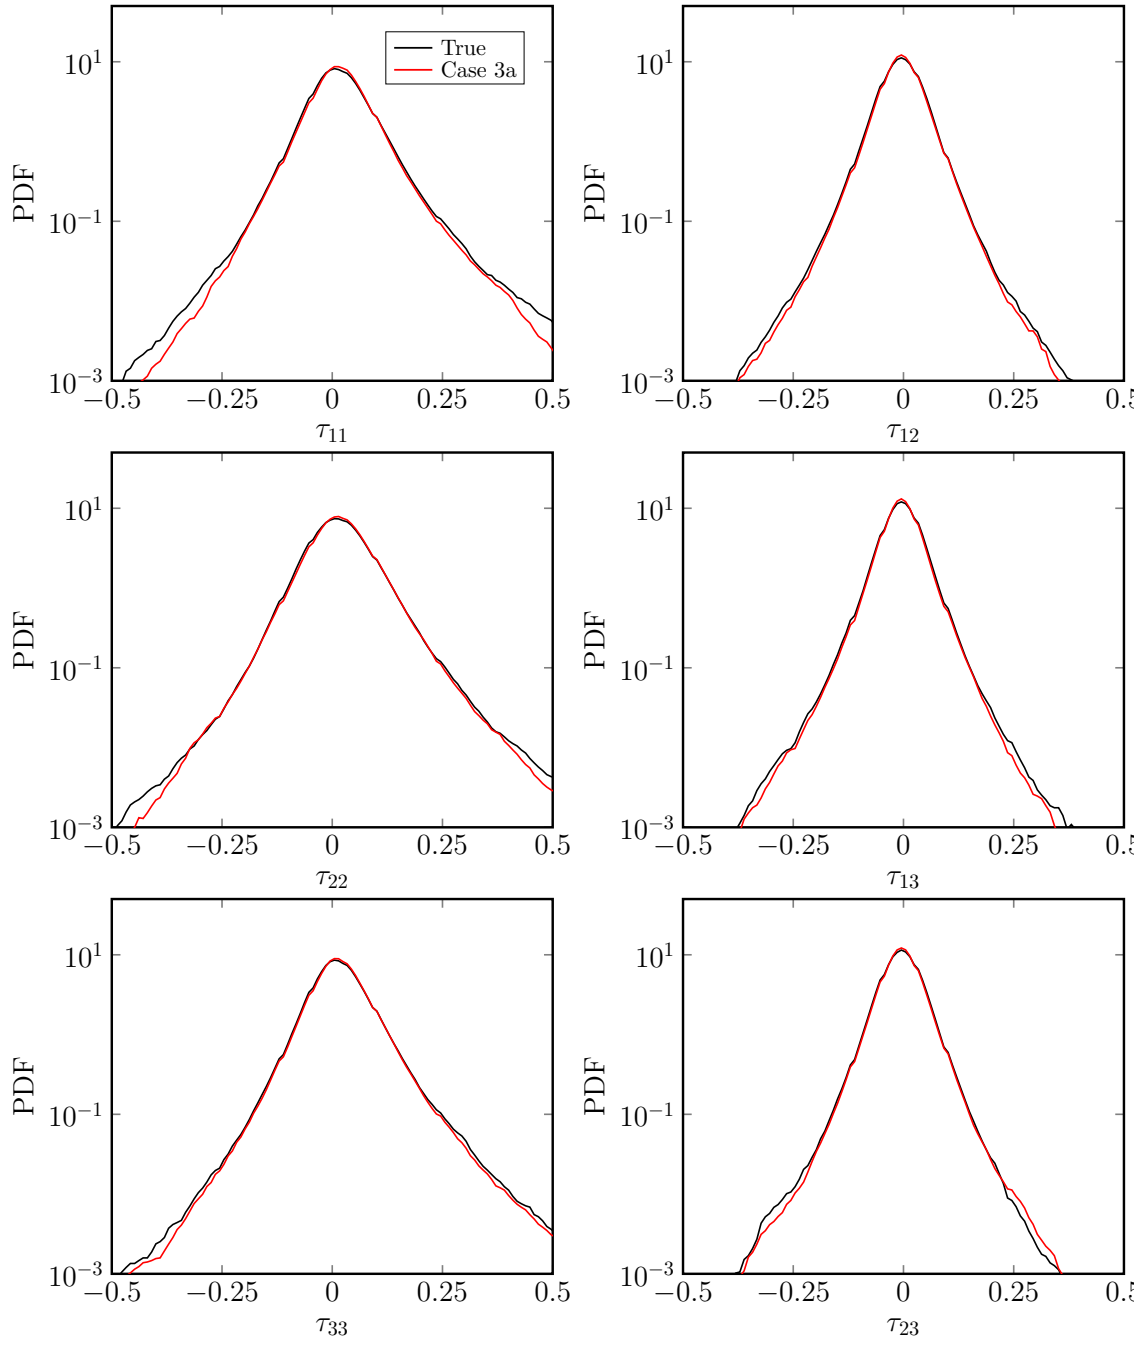


Figure 21. Comparison of probability densities for subgrid stress component fields $\tau_{ij}(\mathbf{x}, t)$ from *a priori* tests of autonomic closure, showing resulting distributions for true subgrid stress fields (black) and from recommended implementation (Case 3a) of autonomic closure (red).




Figure 22. Comparison of probability densities for subgrid production field $P(\mathbf{x}, t)$ from *a priori* tests of autonomic closure, showing results for true subgrid production fields (black) and from recommended implementation (Case 3a) of autonomic closure (red).

tially all resolved scales. This section now compares its performance with that of the dynamic Smagorinsky model [27, 46–59] and the Bardina scale similarity model [43–45]. The former is the most widely used subgrid model for $\tau_{ij}(\mathbf{x}, t)$ in LES, and the latter is the basis for various mixed models [50–57] in which it is typically combined with the Smagorinsky model.

Figure 23 shows typical results comparing the true subgrid stress and production fields $\tau_{ij}(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$ with $\tau_{ij}^F(\mathbf{x}, t)$ and $P^F(\mathbf{x}, t)$ from this implementation of autonomic closure, and with $\tau_{ij}^{DS}(\mathbf{x}, t)$ and $P^{DS}(\mathbf{x}, t)$ from the dynamic Smagorinsky (DS) model and $\tau_{ij}^{BD}(\mathbf{x}, t)$ and $P^{BD}(\mathbf{x}, t)$ from the Bardina scale similarity (BD) model. All results are for the same spectrally sharp LES-scale and test-scale filters and the same test-to-LES filter scale ratio $\alpha \equiv \widehat{\Delta}/\widetilde{\Delta} = 2$ as used throughout. It is apparent in Figure 23 that the results from autonomic closure compare with the true subgrid stress and production fields far better than do the results from either implementation of these two traditional closure models. While the Bardina scale similarity model produces stress fields $\tau_{ij}^{BD}(\mathbf{x}, t)$ that show some of the detailed features in the true $\tau_{ij}(\mathbf{x}, t)$ fields, when contracted with the resolved strain rate tensor in (1.5) the resulting subgrid production field $P^{BD}(\mathbf{x}, t)$ compares relatively poorly with the true pro-

duction field $P(\mathbf{x}, t)$. In comparing subgrid production fields in Figure 23, it is particularly noteworthy that the locations, sizes, and shapes of regions in which large magnitudes of subgrid production are clustered in $P^F(\mathbf{x}, t)$ from autonomic closure agree far better with those in the true $P(\mathbf{x}, t)$ field than do corresponding results from either of the traditional models.

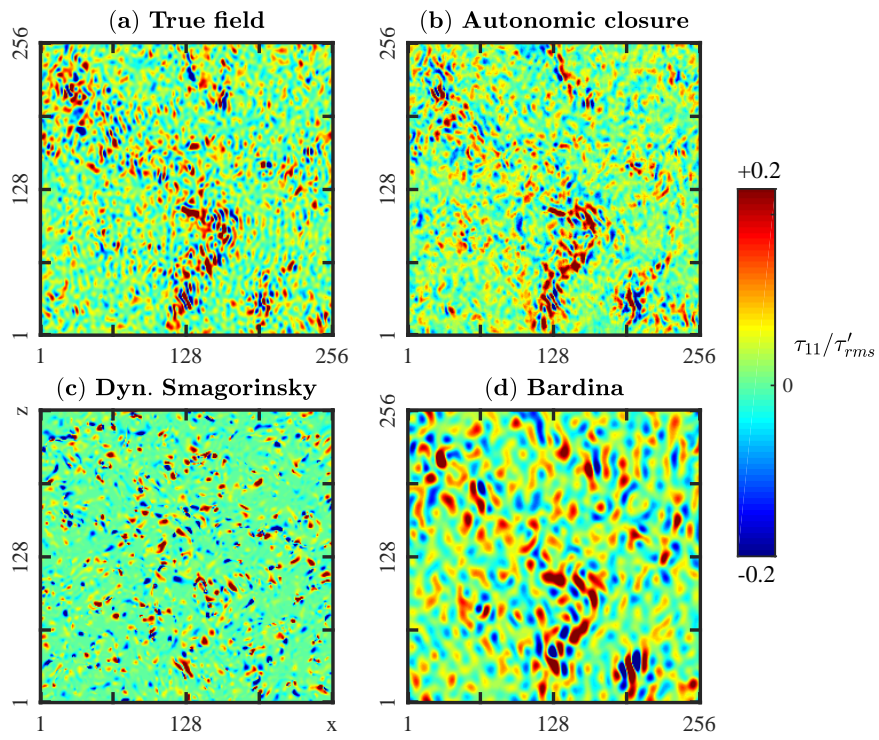


Figure 23 (a). Typical subgrid stress fields $\tau_{11}(\mathbf{x}, t)$.

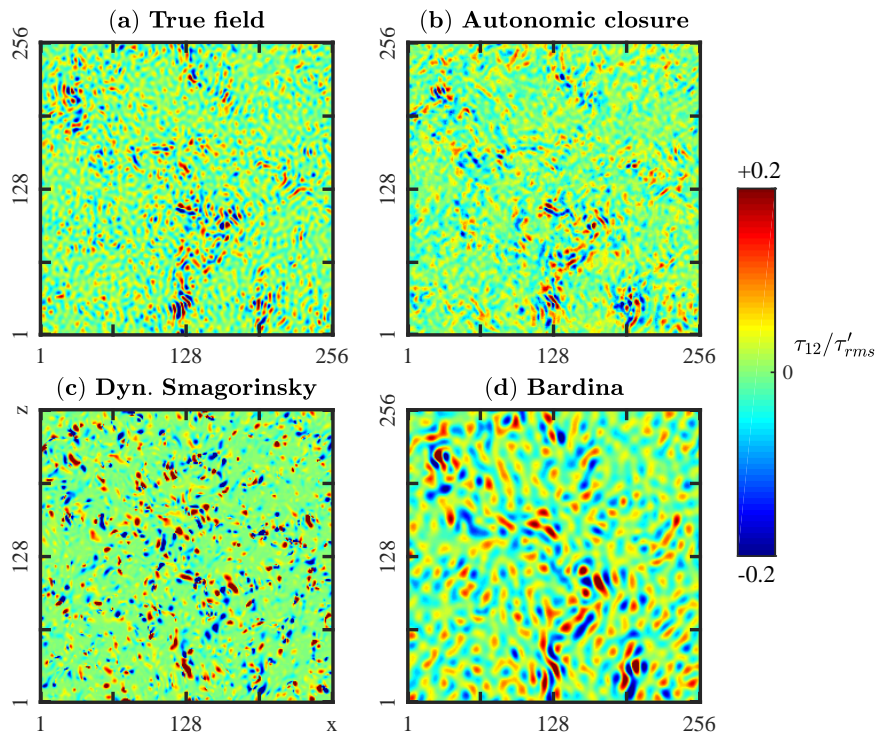


Figure 23 (b). Typical subgrid stress fields $\tau_{12}(\mathbf{x}, t)$.

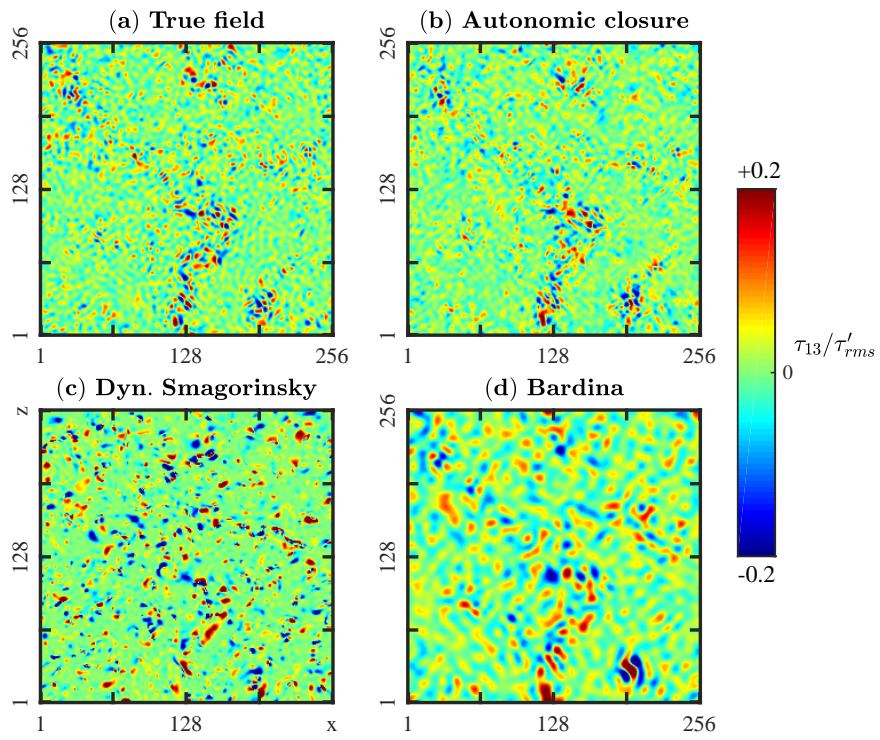


Figure 23 (c). Typical subgrid stress fields $\tau_{13}(\mathbf{x}, t)$.




Figure 23 (d). Typical subgrid stress fields $\tau_{22}(\mathbf{x}, t)$.

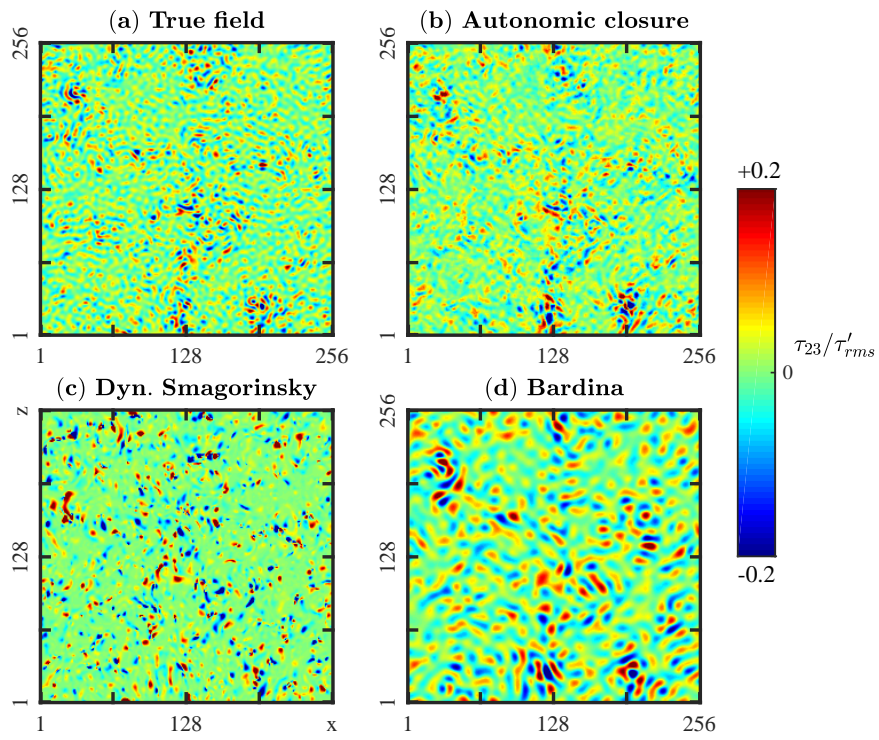


Figure 23 (e). Typical subgrid stress fields $\tau_{23}(\mathbf{x}, t)$.

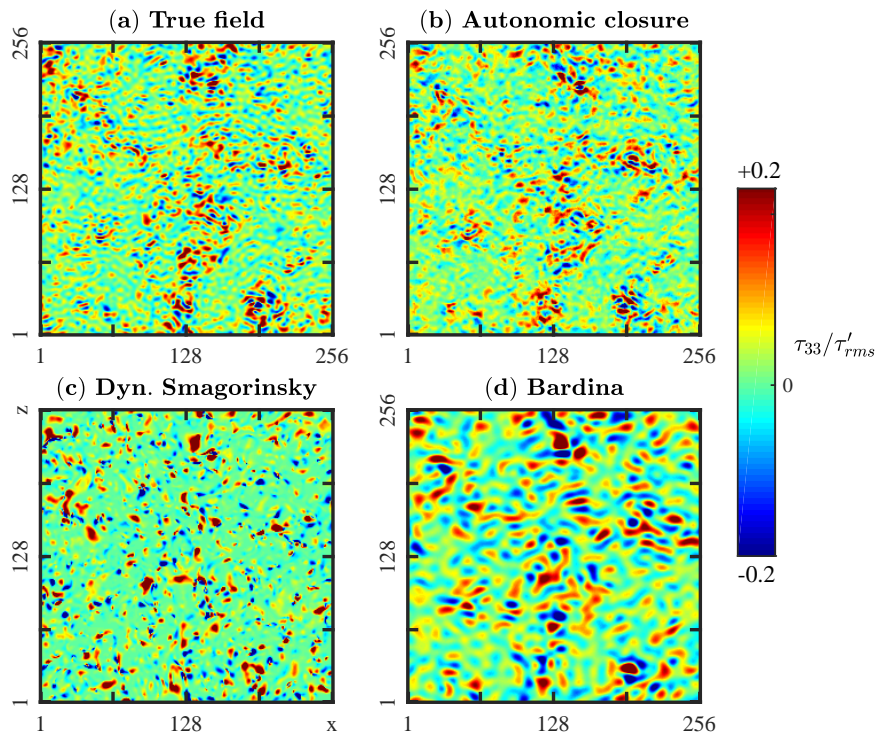


Figure 23 (f). Typical subgrid stress fields $\tau_{33}(\mathbf{x}, t)$.

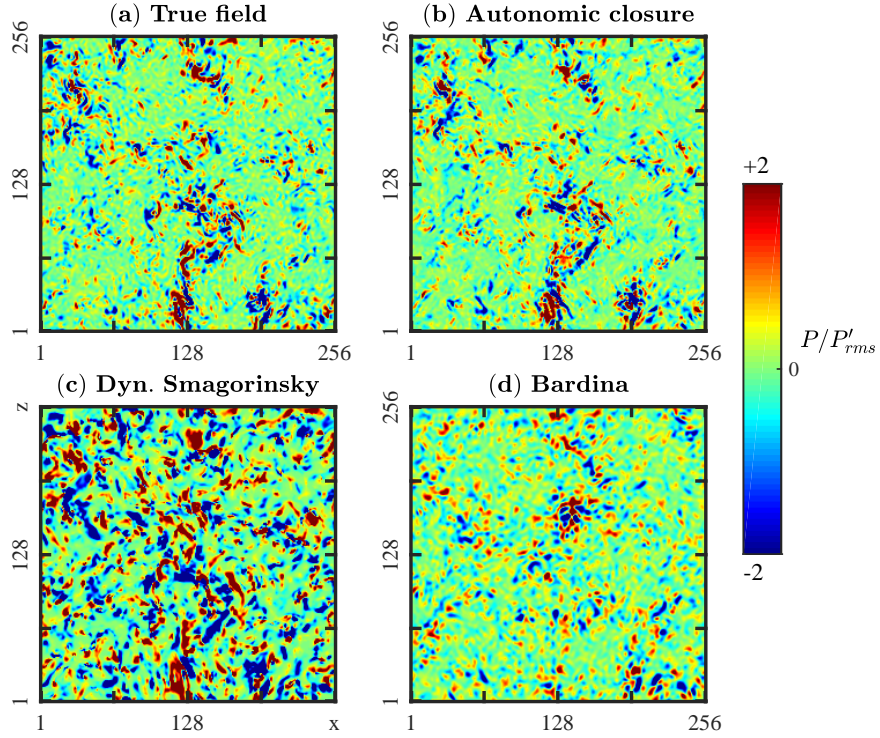


Figure 23 (g). Corresponding subgrid production fields $P(\mathbf{x}, t)$.

Figure 23. Typical comparison of autonomic closure and traditional prescribed closure models, showing (Figure 23a to Figure 23g) typical subgrid stress $\tau_{ij}(\mathbf{x}, t)$ and (Figure 23b) subgrid production $P(\mathbf{x}, t)$; (a) true fields, (b) results from autonomic closure (Case 3a), (c) results from dynamic Smagorinsky (DS) model, and (d) results from Bardina scale-similarity (BD) model; all are for same scale ratio $\alpha \equiv \Delta_\Gamma/\tilde{\Delta} = 2$.

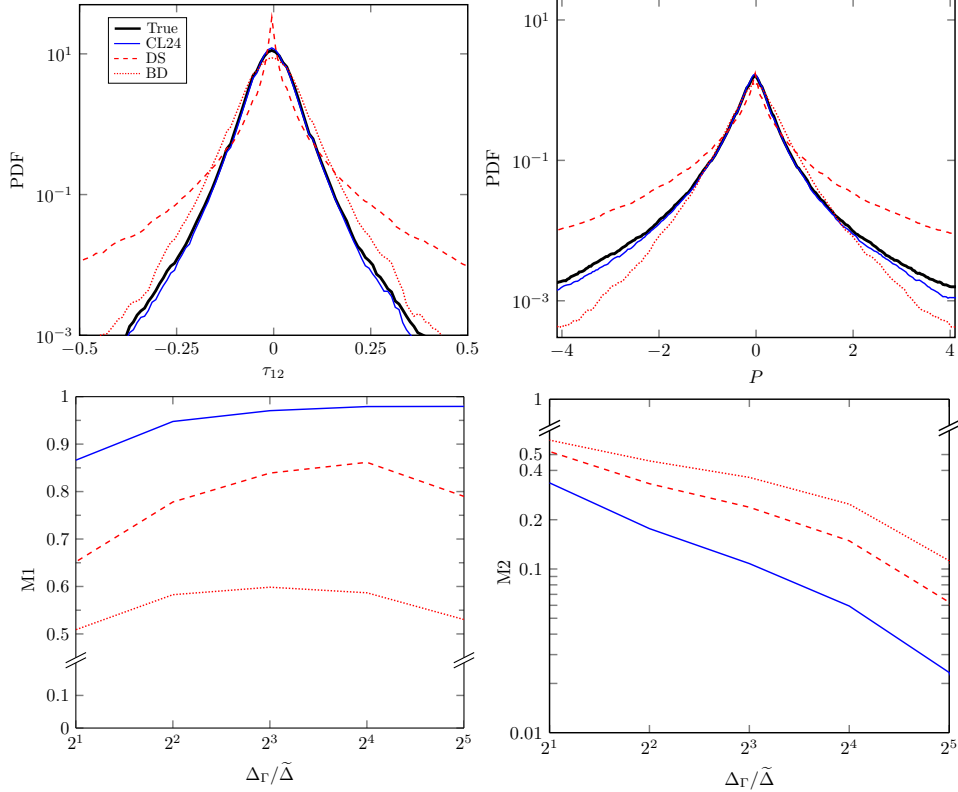


Figure 24. Typical comparison of autonomic closure and traditional prescribed closure models, showing (top) pdfs of typical subgrid stress component τ_{ij} and subgrid production P , and (bottom) M_1 and M_2 variation with scale ratio $\Delta_\Gamma/\tilde{\Delta}$ for subgrid production support-density fields; true fields (black), autonomic closure (blue), dynamic Smagorinsky model (red dashed), Bardina scale-similarity model (red dotted); all are for same scale ratio $\alpha \equiv \Delta_\Gamma/\tilde{\Delta} = 2$.

Figure 24a,b show comparisons of resulting probability densities of the subgrid stress and subgrid production from each of these subgrid closures. It is apparent that the statistical distributions from autonomic closure match the true distributions very closely. Additionally, Figure 24c,d show the support-density metrics M_1 and M_2 comparing the spatial support on which large positive and negative values of subgrid production are concentrated in Figure 23. It is apparent in Figure 24, and consistent with the visual comparisons in Figure 23, that the spatial structure of the support-density on which large values of positive and negative subgrid production values are concentrated in $P^F(\mathbf{x}, t)$ from autonomic closure agrees far better with the true

subgrid production field $P(\mathbf{x}, t)$ than do the corresponding results $P^{DS}(\mathbf{x}, t)$ and $P^{BD}(\mathbf{x}, t)$ from the dynamic Smagorinsky and Bardina scale similarity models.

Note in Table 1 that the average subgrid production $\langle P^F \rangle$ from autonomic closure closely matches the true value ε , corresponding to average energy transfer from the resolved scales into the subgrid scales at a rate just slightly higher than ε . By comparison, $\langle P^{BD} \rangle$ from the Bardina scale similarity model is seen in Table 1 to have the opposite sign, and thus on average transfers energy into the resolved scales. This is consistent with the widely reported observation that the scale similarity model leads to unstable simulations unless it is coupled with an added dissipative model to ensure net average energy transfer out of the resolved scales. Table 1 shows that $\langle P^{DS} \rangle$ from the dynamic Smagorinsky model has the correct sign and thus on average transfers energy out of the resolved scales. However, Figure 23 shows that $P^{DS}(\mathbf{x}, t)$ compares very poorly to the true $P(\mathbf{x}, t)$ field, with large positive and negative values of subgrid production being highly overrepresented, and with regions in which large forward and backward scatter are seen in $P^{DS}(\mathbf{x}, t)$ poorly matching those in the true field $P(\mathbf{x}, t)$. These factors likely contribute to the widely reported need for limiters, added dissipation, or other *ad hoc* treatments to keep the dynamic Smagorinsky model computationally stable.

In contrast, autonomic closure produces $\tau_{ij}^F(\mathbf{x}, t)$ and $P^F(\mathbf{x}, t)$ in Figures 12 and 23 that closely match even detailed features in the true fields $\tau_{ij}(\mathbf{x}, t)$ and $P(\mathbf{x}, t)$, and thus also closely match the stress and production statistics in Figures 21, 22 and 24, and lead to $\langle P^F \rangle$ that closely matches the true value ε in Table 1. Collectively, these factors suggest it may be possible for this Case 3a implementation of autonomic closure to be stable without the need for limiters, added dissipation, or other *ad hoc* treatments when implemented in large eddy simulations, though this can only be assessed via future *aposteriori* tests.

Note also in Table 1 that the computational time for this Case 3a implementation of autonomic closure is only about an order of magnitude larger than that needed to evaluate the dynamic Smagorinsky model or the Bardina scale similarity model. Since the computational time for subgrid stress evaluation is typically only a small fraction of the total computational

time needed for LES, the Case 3a implementation of autonomic closure should be efficient enough for use in practical large eddy simulations.

INVARIANT REPRESENTATIONS IN AUTONOMIC CLOSURE

In Chapters 2-4, autonomic closure was implemented using a local series representation of the subgrid terms in the values of the primitive variables of a large eddy simulation. For the subgrid stress these are the values of the velocity components and pressure at the 27 points on the $3 \times 3 \times 3$ stencil. Translational invariance was enforced by subtracting the stencil center-point velocity from the velocities at all the stencil points, thus in effect retaining only velocity gradients on the stencil. To keep the number of terms in this series manageable, the series was truncated after second-order terms. Furthermore, keeping pressure values in this series was shown to provide negligible benefit when second-order terms are retained in the series. This led to a second-order velocity-only representation as a 244-term series for each of the six unique components of the subgrid stress. Autonomic closure then determines the optimal local coefficients for each of these 6×244 terms.

Although this truncated series representation for τ_{ij} was found in Chapter 4 to provide remarkably accurate results, the truncation is undeniably arbitrary and likely to be less computationally efficient than a more “appropriate” representation based on a frame-invariant form for τ_{ij} . Such an “appropriate” representation is one involving only terms that preserve the tensor invariance properties of τ_{ij} , specifically the rank, symmetry, rotation, and reflection invariance of rank-2 symmetric tensors. This chapter develops such representations for τ_{ij} , which then are evaluated in Chapter 6 to identify the most accurate and efficient implementation for autonomic closure.

5.1 Frame-Invariant Representation Theory

Representation theory formally provides the most general representation of any given quantity, such as the subgrid stress tensor τ_{ij} , in terms of any set of other quantities, such as the velocities \mathbf{u}_m and pressures p_m on the stencil, that preserves a given set of imposed constraints that

the representation must satisfy. With regard to the generalized representation F_{ij} for the sub-grid stress τ_{ij} in autonomic closure, these constraints consist of the rank, symmetry, rotation, and reflection invariance of rank-2 symmetric tensors. Translational invariance is enforced separately as described at the beginning of this chapter. Symmetry simply requires $F_{ij} = F_{ji}$. However, invariance under frame rotation and reflection places a far larger set of less obvious constraints on the representation.

With regard to frame rotation and reflection invariance, for any proper orthogonal tensor \mathbf{Q} (i.e., $Q_{ik}Q_{kj} = \delta_{ij}$ with $\det \mathbf{Q} = 1$) that rotates Cartesian coordinate frame \mathbf{x} into a new Cartesian frame \mathbf{x}' as $x_i = Q_{ij}x_j$, the stress tensors in the two frames must be related as $\tau'_{ij} = Q_{ik} \tau_{kl} Q_{lj}^T$, so that the scalar invariants $I = \tau_{ii}$, $II = \tau_{ij}\tau_{ji}$, and $III = \tau_{ij}\tau_{jk}\tau_{ki}$ are the same in both coordinate frames. The stress tensor is then said to be “frame invariant”; specifically, in an n -dimensional Cartesian space it satisfies the symmetry, rotation, and reflection properties associated with the special orthogonal group $SO(n)$. Any tensorally-correct representation F_{ij} for τ_{ij} must preserve these symmetry, rotation, and reflection properties.

In Chapters 2-4 these frame-invariance properties were only indirectly (and weakly) incorporated in autonomic closure, by training the solution for the series coefficients on test-scale stresses that inherently satisfy these same invariance properties. In principle, these invariance properties could be directly enforced by using an alternative representation for F_{ij} in which each term satisfies these frame-invariance properties. As shown in the following sections, such an approach can lead to a finite series representation of τ_{ij} consisting of a smaller number of terms that does not require truncation. Moreover, this finite series representation can be “complete” in the sense that it contains all possible tensorally-correct combinations of an assumed set of quantities that τ_{ij} can depend on. Such a complete frame-invariant representation would presumably be the most accurate and efficient formulation of autonomic closure.

5.2 Invariant Representations of τ_{ij} in the Stencil Velocities \mathbf{u}_m

Smith (1971) provides a general formulation for the *complete* tensor basis needed to represent any symmetric rank-two tensor in terms of any number of vectors and rank-two tensors.

Further, Pennisi and Trovato (1987) have shown that the tensor basis set resulting from Smith’s general formulation is *minimal*, meaning there is no smaller tensor basis set that can give a complete frame-invariant representation of the symmetric rank-two tensor.

It is thus possible from Smith’s general formulation to obtain a complete and minimal tensor polynomial representation for τ in terms of the velocities \mathbf{u}_m on the $3 \times 3 \times 3$ stencil. This section uses this approach to determine a frame-invariant representation of the symmetric rank-two stress tensor τ in terms of the $m = 1, 2, \dots, P = 27$ stencil-point velocity vectors \mathbf{u}_m of the form $\tau_{ij} = F_{ij}(\mathbf{u}_m)$, meaning a representation that preserves the rank, symmetry, rotation and reflection properties of τ_{ij} .

Following Smith (1971; Eq. 4.5), any such rank-two symmetric tensor τ can be represented in a complete and minimal tensor polynomial in terms of any number P of vectors \mathbf{u}_m as

$$\tau = c_0 \mathbf{I} + \sum_{m=1}^P c_m (\mathbf{u}_m \otimes \mathbf{u}_m) + \sum_{m=1}^P \sum_{n=m+1}^P c_{mn} (\mathbf{u}_m \otimes \mathbf{u}_n + \mathbf{u}_n \otimes \mathbf{u}_m), \quad (5.1)$$

where \otimes denotes the “dyadic product” (or “outer product”, or “tensor product”) defined for any two vectors $\mathbf{a} = (a_1, a_2, a_3)$ and $\mathbf{b} = (b_1, b_2, b_3)$ as

$$\mathbf{a} \otimes \mathbf{b} = \begin{pmatrix} a_1 b_1 & a_1 b_2 & a_1 b_3 \\ a_2 b_1 & a_2 b_2 & a_2 b_3 \\ a_3 b_1 & a_3 b_2 & a_3 b_3 \end{pmatrix}. \quad (5.2)$$

Unlike in the earlier series representations in Chapters 2-4, note that all six independent components of τ in (5.1) have the *same* set of coefficients c , so that (5.1) preserves all invariants of τ in *any* coordinate frame. In contrast, the earlier *ad hoc* series representations were not invariant-preserving and had separate coefficients for each ij -component of τ_{ij} .

Since in (5.1) the \mathbf{u}_m are the velocities at each of the 27 stencil points, the symmetric dyadic products in that representation are simply all possible frame-invariance-preserving combinations of all second-order products (colocated and non-colocated) of the velocity components. The first sum gives the 27 colocated velocity products, which include both square ($i = j$) and non-square ($i \neq j$) terms. The second sum gives the $(27^2 - 27)/2 = 351$ non-colocated

velocity products; note how the dyadic products in this term are constructed to preserve the ij -symmetry of τ .

Strictly speaking, representation theory dictates that the scalar-valued coefficients c in (5.1) can, at most, be functions of all scalar invariants that can be formed from the set of vectors \mathbf{u}_m . From Smith (1971; Eq. 2.41), for a representation involving only the vectors \mathbf{u}_m , the only such scalar invariants are $I_{\alpha\alpha} \equiv \mathbf{u}_\alpha \cdot \mathbf{u}_\alpha$ and $I_{\alpha\beta} \equiv \mathbf{u}_\alpha \cdot \mathbf{u}_\beta$ where $\alpha = 1, 2, \dots, P$ and $\beta > \alpha$. In principle, one could try to somehow discover how each of the c 's in (5.1) depend on the P scalar invariants $I_{\alpha\alpha}$ and the $(P^2 - P)/2$ scalar invariants $I_{\alpha\beta}$, and if successful in doing so would have arrived at the only possible tensorally-correct representation of the form $\tau_{ij} = F_{ij}(\mathbf{u}_m)$. However, the whole point of autonomic closure is precisely to avoid a need to propose models for coefficients $c = f(I_{11}, \dots, I_{PP}, I_{12}, \dots, I_{P-1P})$ in representations for τ , and to instead use tools of optimization and machine learning to autonomically discover the best values for these coefficients at each point and time in a simulation.

Returning to (5.2), it might be tempting to discard the term $c_0\mathbf{I}$ by arguing that the definition of τ , namely

$$\tau = \widetilde{\mathbf{u} \otimes \mathbf{u}} - \widetilde{\mathbf{u}} \otimes \widetilde{\mathbf{u}} \quad \text{or equivalently} \quad \tau_{ij} \equiv \widetilde{u_i u_j} - \widetilde{u}_i \widetilde{u}_j, \quad (5.3)$$

requires $\tau = 0$ when $\mathbf{u} \equiv 0$. One might thus consider forcing $c_0 \equiv 0$ to ensure (5.1) satisfies this requirement. However, recall from above that c_0 is in principle a scalar-valued function f of the scalar invariants $\mathbf{u}_\alpha \cdot \mathbf{u}_\alpha$ and $\mathbf{u}_\alpha \cdot \mathbf{u}_\beta$, so when all the $\mathbf{u}_{\alpha,\beta} \equiv 0$ then these scalar invariants are all zero, in which case it is possible (even likely) that $c_0 = f(0, \dots, 0) = 0$. This then allows $c_0 = 0$ when $\mathbf{u} \equiv 0$ and thereby satisfies $\tau = 0$ when $\mathbf{u} \equiv 0$ even when the $c_0\mathbf{I}$ term in (5.1) is retained. Furthermore, when $\mathbf{u} \neq 0$ we may expect $c_0 \neq 0$, in which case again the $c_0\mathbf{I}$ term in (5.1) must be retained. Thus in all cases, the $c_0\mathbf{I}$ term in (5.1) must be retained. Only when τ in (5.1) is the deviatoric stress should $c_0 \equiv 0$.

Note that, since the c 's in (5.1) are the *same* for all six independent components of τ , there are only $1 + 27 + 351 = 379$ coefficients involved in the entire representation of τ , even though (5.1) includes all possible invariance-preserving non-located velocity products. The original

series was not invariance-preserving, and required $6 \times 3403 = 20,418$ coefficients to represent the entire tensor τ . If only colocated velocity products were retained in the original series it required $6 \times 244 = 1464$ coefficients, whereas in (5.1) including only colocated velocity products amounts to omitting the final term, which leaves just $1 + 27 = 28$ coefficients that must be found to completely represent τ .

The generalized representation in (5.1) has several advantages over the original truncated *ad hoc* series representations in Chapters 2-4, including:

- Unlike the *ad hoc* representations, (5.1) is tensorally correct in that it represents the *entire* tensor τ as a single polynomial in which each term is a rank-two tensor that preserves the symmetry, rotation and reflection properties of τ .
- While the *ad hoc* series representation required truncation after second-order terms to be practically implementable, (5.1) is a *complete* representation of τ in terms of the stencil velocities \mathbf{u}_m . It involves no truncation; representation theory limits it terms no higher than second-order in the velocities \mathbf{u}_m .
- Galilean (translation) invariance is satisfied by (5.1) when the stencil velocities \mathbf{u}_m are understood to have the stencil center-point velocity subtracted from them, so that the \mathbf{u}_m denote the resulting velocity differences. This also causes the resulting \mathbf{u}_m at the stencil center point to be zero, thereby reducing P from 27 to 26 independent velocity vectors. That in turn reduces the number of coefficients in (5.1) to $1 + 26 + (26^2 - 26)/2 = 352$, and if only colocated velocity products are included then the number of coefficients is only $1 + 26 = 27$.
- The representation in (5.1) involves a far smaller number N of degrees-of-freedom (i.e., coefficients c_i) than did the original series, and thus would allow *far* faster “on-the-fly” implementations of autonomic closure in an LES code. Specifically, it involves a factor of $20,418/379 = 54$ fewer degrees-of-freedom N in its representation of τ than did the truncated series in the components of \mathbf{u}_m . Since the computational burden of autonomic

- closure scales as $N^3[1 + M/N]$, with $N > M$ this reduction in N should reduce the computational time by a factor of $54^3 = 150,000$ relative to the original truncated series.
- From the reduced number of degrees-of-freedom N needed in (5.1) to represent τ , far fewer training points should be needed to determine these coefficients while keeping the relative training ratio N/M the same as in Chapter 4. Beyond the resulting reduction in computational cost, this should allow an even smaller bounding box than was needed for the Case 3a implementation from Chapter 4, and thus should allow a more local (and thereby a more accurate) closure.
 - Beyond the reduction in N , further gains in efficiency result from the fact that a frame-invariant representation as in (5.1) involves only a *single* set of coefficients, whereas the original *ad hoc* series representation required a separate set of coefficients for each components of τ_{ij} . Specifically, the *ad hoc* series representation required $6 \times N$ coefficients to determine τ_{ij} , whereas a frame-invariant representation requires only $1 \times N$ coefficients.
 - Since a frame-invariant representation reduces the number of coefficients from $6 \times N$ to $1 \times N$, each of the M training points provides *six* sets of training data, one from each of the six independent components of T_{ij} . Thus the M training point locations provide a total of $M' = 6M$ training data sets. The minimum required training data per degree-of-freedom is then $M'/N \approx \mathcal{O}(4)$, allowing fewer training points and thereby smaller bounding boxes to be used, which provide further increases in efficiency and accuracy.

5.3 Invariant Representations of τ_{ij} in \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$, and $\nabla\mathbf{R}$

This section uses the general formulation of Smith (1971) to obtain a complete and minimal representation of the symmetric subgrid stress tensor τ_{ij} in terms of the strain and rotation rate tensors \mathbf{S} and \mathbf{R} and rank-two products up to second order of their gradients $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$. In particular, Smith's framework guarantees that the resulting tensor polynomial representation will satisfy the rank, symmetry, rotation, and reflection properties of τ_{ij} . Moreover, Smith's

formulation was shown by Pennisi et al. (1987) to be minimal, thus there is no smaller number of symmetric basis tensors that can form a complete polynomial representation of τ_{ij} .

Unlike other invariance-preserving general representations for the subgrid stress, which typically assume that τ_{ij} depends only on the strain rate tensor $\mathbf{S} \equiv S_{ij}$, or only on the strain rate and rotation rate tensors \mathbf{S} and $\mathbf{R} \equiv R_{ij}$, in autonomic closure the relative velocities on a $3 \times 3 \times 3$ stencil provide not only the strain and rotation rate tensors \mathbf{S} and \mathbf{R} but also their gradients $\nabla\mathbf{S} \equiv \partial S_{ij}/\partial x_k$ and $\nabla\mathbf{R} \equiv \partial R_{ij}/\partial x_k$. This allows an even far more general tensor invariant formulation for τ_{ij} than has previously been obtained [72, 83, 85, 86].

The resulting complete and minimal representation will be seen to consist of 1570 tensor terms that each involve products up to *fourth* order in the underlying rank-two tensors \mathbf{M}_i and \mathbf{W}_p , each of which in turn involve products up to second order in \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$. Since the strain and rotation rate tensors \mathbf{S} and \mathbf{R} , and gradients $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$, are linear in the components of the velocities \mathbf{u}_m at the $m = 1, 2, \dots, P = 27$ points on the local $3 \times 3 \times 3$ stencil, each term in this frame-invariant representation ultimately involves products up to *eighth* order in these velocity components. Based on the observation in Chapter 4 that velocity products no higher than second order are needed to obtain accurate representations for τ_{ij} , the frame-invariant representation developed in this section will also be truncated at this order,

5.3.1 Invariance-Preserving Tensor Representations of τ_{ij}

5.3.1.1 Linear representations in \mathbf{S}

The most widely used models assume that the turbulent stress depends only linearly on the resolved strain rate tensor S_{ij} and the identity tensor δ_{ij} . There are only two independent tensors that can be formed from S_{ij} and δ_{ij} that are linear in S_{ij} and preserve the rank, symmetry, rotation, and reflection properties of τ_{ij} , namely

$$\mathbf{m}^{(0)} = S_{kk}\delta_{ij} \tag{5.4a}$$

$$\mathbf{m}^{(1)} = S_{ij} \tag{5.4b}$$

The most general representation for the subgrid stress must then be a linear combination of these two invariance-preserving tensors $\mathbf{m}^{(\alpha)}$ where $\alpha = 0, 1$, and thus can be written as

$$\tau_{ij} = \sum_{\alpha=0}^1 c_{\alpha} \mathbf{m}^{(\alpha)} = c_0 S_{kk} \delta_{ij} + c_1 S_{ij}, \quad (5.5)$$

where the coefficients c_0 and c_1 are scalars and so can be functions only of the four invariance-preserving scalars that can be formed from \tilde{S}_{ij} and δ_{ij} , namely $c_{\alpha} = f_{\alpha}(\mathbf{I}_0, \mathbf{I}_1, \mathbf{I}_2, \mathbf{I}_3)$ where

$$\mathbf{I}_0 = tr(\mathbf{I}) = \delta_{ii} \quad (5.6a)$$

$$\mathbf{I}_1 = tr(\mathbf{S}) = S_{ii} \quad (5.6b)$$

$$\mathbf{I}_2 = tr(\mathbf{S}^2) = S_{ij} S_{ij} \quad (5.6c)$$

$$\mathbf{I}_3 = tr(\mathbf{S}^3) = S_{ij} S_{jk} S_{ki}. \quad (5.6d)$$

Equation (5.5) can be equivalently written for the deviatoric stress $\tau_{ij}^{dev} \equiv (\tau_{ij} - \tau_{kk} \delta_{ij}/3)$ and the deviatoric strain rate $S_{ij}^{dev} \equiv (S_{ij} - S_{kk} \delta_{ij}/3)$, for which $S_{kk}^{dev} \equiv 0$. In that case $\mathbf{m}^{(0)} \equiv 0$, so analogous to (5.5) the most general representation is

$$\tau_{ij}^{dev} = \sum_{\alpha=0}^1 c_{\alpha} \mathbf{m}^{(\alpha)} = c_1 S_{ij}^{dev}, \quad (5.7)$$

where c_1 (generally called the ‘‘subgrid viscosity’’ and denoted ν_{sgs}) can depend only on the four scalar invariants $\mathbf{I}_0, \mathbf{I}_1, \mathbf{I}_2, \mathbf{I}_3$ in (5.6). The invariance-preserving form for τ_{ij} in (5.5), or equivalently in (5.7), is the basis of so-called *linear* subgrid-scale models, since they assume the subgrid stress to be linearly related to the strain rate tensor, with various linear models differing only in their choice of how the coefficient ν_{sgs} depends on the scalar invariants $\mathbf{I}_0, \mathbf{I}_1, \mathbf{I}_2, \mathbf{I}_3$. However, for the purposes of this study the equivalent form in (5.5) is more useful, since it shows how this most general *linear* representation of τ_{ij} can be obtained from tensor invariance, under the assumption that τ_{ij} depends only on S_{ij} and δ_{ij} , based on the corresponding tensor bases $\mathbf{m}^{(\alpha)}$.

5.3.1.2 Nonlinear representations in \mathbf{S} and \mathbf{R}

Similarly, there are ‘‘nonlinear’’ models that allow τ_{ij} to depend not only on the strain rate tensor S_{ij} and the identity tensor δ_{ij} , but also on the rotation rate tensor R_{ij} . Such models

represent τ_{ij} in a complete set of tensor bases $\mathbf{m}^{(\alpha)}$ obtained from combinations of S_{ij} , R_{ij} , and δ_{ij} that preserve the rank, symmetry, translation, and rotation properties of τ_{ij} . For instance, Lumley (1970), Pope (1974), Lund & Novikov (1992), and Gatski & Speziale (1993) propose an 11-element set of invariance-preserving tensors that can be formed from S_{ij} , R_{ij} , and δ_{ij} , namely

$$\mathbf{m}^{(0)} = \mathbf{I} = \delta_{ij} \quad (5.8a)$$

$$\mathbf{m}^{(1)} = \mathbf{S} = \tilde{S}_{ij} \quad (5.8b)$$

$$\mathbf{m}^{(2)} = \mathbf{S}^2 = \tilde{S}_{ik}\tilde{S}_{kj} \quad (5.8c)$$

$$\mathbf{m}^{(3)} = \mathbf{R}^2 = \tilde{R}_{ik}\tilde{R}_{kj} \quad (5.8d)$$

$$\mathbf{m}^{(4)} = \mathbf{SR} - \mathbf{RS} = \tilde{S}_{ik}\tilde{R}_{kj} - \tilde{R}_{ik}\tilde{S}_{kj} \quad (5.8e)$$

$$\mathbf{m}^{(5)} = \mathbf{S}^2\mathbf{R} - \mathbf{RS}^2 = \tilde{S}_{ik}\tilde{S}_{kl}\tilde{R}_{lj} - \tilde{R}_{ik}\tilde{S}_{kl}\tilde{S}_{lj} \quad (5.8f)$$

$$\mathbf{m}^{(6)} = \mathbf{SR}^2 + \mathbf{R}^2\mathbf{S} = \tilde{S}_{ik}\tilde{S}_{kl}\tilde{R}_{lj} + \tilde{R}_{ik}\tilde{S}_{kl}\tilde{S}_{lj} \quad (5.8g)$$

$$\mathbf{m}^{(7)} = \mathbf{S}^2\mathbf{R}^2 + \mathbf{R}^2\mathbf{S}^2 = \tilde{S}_{ik}\tilde{S}_{kl}\tilde{R}_{lm}\tilde{R}_{mj} + \tilde{R}_{ik}\tilde{R}_{kl}\tilde{S}_{lm}\tilde{S}_{mj} \quad (5.8h)$$

$$\mathbf{m}^{(8)} = \mathbf{SRS}^2 - \mathbf{S}^2\mathbf{RS} = \tilde{S}_{ik}\tilde{R}_{kl}\tilde{S}_{lm}\tilde{S}_{mj} - \tilde{S}_{ik}\tilde{S}_{kl}\tilde{R}_{lm}\tilde{S}_{mj} \quad (5.8i)$$

$$\mathbf{m}^{(9)} = \mathbf{RSR}^2 - \mathbf{R}^2\mathbf{SR} = \tilde{R}_{ik}\tilde{S}_{kl}\tilde{R}_{lm}\tilde{R}_{mj} - \tilde{R}_{ik}\tilde{R}_{kl}\tilde{S}_{lm}\tilde{R}_{mj} \quad (5.8j)$$

$$\mathbf{m}^{(10)} = \mathbf{RS}^2\mathbf{R}^2 - \mathbf{R}^2\mathbf{S}^2\mathbf{R} = \tilde{R}_{ik}\tilde{S}_{kl}\tilde{S}_{lm}\tilde{R}_{mn}\tilde{R}_{nj} - \tilde{R}_{ik}\tilde{R}_{kl}\tilde{S}_{lm}\tilde{S}_{mn}\tilde{R}_{nj} \quad (5.8k)$$

Note that (5.8a-k) involve no matrix powers higher than two; i.e., \mathbf{S}^2 and \mathbf{R}^2 appear, but neither \mathbf{S}^3 nor \mathbf{R}^3 nor any higher-order matrix powers appear. This is a result of the Cayley-Hamilton theorem for any square matrix \mathbf{A} , which relates its matrix powers of order three (\mathbf{A}^3) and higher to linear combinations of lower matrix powers. It is this fact that limits the complete set of such invariance-preserving tensors $\mathbf{m}^{(\alpha)}$ to the *finite* number of *independent* tensors in (5.8).

Thus, if it is assumed that τ_{ij} depends only on S_{ij} , R_{ij} , and δ_{ij} , then analogous to (5.5) the most general representation must be a linear sum of these invariance-preserving tensors,

namely

$$\tau_{ij} = \sum_{\alpha=0}^{10} c_{\alpha} \mathbf{m}^{(\alpha)}, \quad (5.9)$$

where the coefficients $c_{\alpha} = f(\mathbf{I}_0, \dots, \mathbf{I}_6)$ can only be functions of the seven unique independent invariance-preserving scalar invariants that can be formed from S_{ij} , R_{ij} , and δ_{ij} , namely

$$\mathbf{I}_0 = tr(\mathbf{I}) = \delta_{ii} \quad (5.10a)$$

$$\mathbf{I}_1 = tr(\mathbf{S}^2) = S_{ij}S_{ij} \quad (5.10b)$$

$$\mathbf{I}_2 = tr(\mathbf{R}^2) = R_{ij}R_{ij} \quad (5.10c)$$

$$\mathbf{I}_3 = tr(\mathbf{S}^3) = S_{ik}S_{kl}S_{li} \quad (5.10d)$$

$$\mathbf{I}_4 = tr(\mathbf{SR}^2) = S_{ik}R_{kl}R_{li} \quad (5.10e)$$

$$\mathbf{I}_5 = tr(\mathbf{S}^2\mathbf{R}^2) = S_{ik}S_{kl}R_{lm}R_{mi} \quad (5.10f)$$

$$\mathbf{I}_6 = tr(\mathbf{S}^2\mathbf{R}^2\mathbf{SR}) = S_{ik}S_{kl}R_{lm}R_{mn}S_{no}R_{oi} \quad (5.10g)$$

Equation (5.9), with (5.8) and (5.10), is a complete general tensor polynomial in S_{ij} , R_{ij} , and δ_{ii} that preserves the rank, symmetry, rotation, and reflection properties of τ_{ij} . Under the assumption that τ_{ij} depends only on \mathbf{S} , \mathbf{R} and \mathbf{I} , all valid representations of the stress tensor τ_{ij} must be expressible as in (5.9).

Lund & Novikov (1992) further showed that, under the additional assumption that the strain rate \mathbf{S} is not in an axisymmetric state (i.e., when \mathbf{S} does not have repeated eigenvalues) and when the vorticity vector obtained from \mathbf{R} is not aligned with any of the strain rate eigenvectors, then the number of *independent* $\mathbf{m}^{(\alpha)}$ in (5.9) is reduced from eleven to six, and the number of independent scalar invariants in (5.10a-g) is reduced from seven to five.

5.3.2 Toward the Most-General Frame-Invariant Representation

In principle, the subgrid stress in autonomic closure can be represented by the invariance-preserving series in (5.9) in terms of the tensor bases $\mathbf{m}^{(\alpha)}$ in (5.8). The coefficients c_{α} would then be determined locally via the autonomic closure methodology, rather than from a prescribed model that relates the coefficients to the seven scalar invariants I_{α} in (5.9). Doing this

would allow implementing autonomic closure with a representation having far fewer coefficients than the 244-term series representation in the primitive variables identified in Chapter 4, which should be more computationally efficient (though evaluating the $\mathbf{m}^{(\alpha)}$ to build the $\widehat{\mathbf{V}}$ matrix may consume part of the reduction in computational time). More important, such a formulation will no longer be based on an arbitrarily truncated series, but instead will be based on the most general *complete* series representation (under the assumption that τ_{ij} depends only on S_{ij} , R_{ij} , and δ_{ij}) that is consistent with the tensor properties of the subgrid stress.

Note the representation in (5.9) with $\mathbf{m}^{(\alpha)}$ from (5.8) involves products of velocity components up to 5th order (via $\mathbf{m}^{(10)}$). That representation in $\mathbf{m}^{(\alpha)}$ is the most general tensor polynomial that preserves all rank, symmetry, rotation, and reflection invariance properties of τ_{ij} under the assumption that τ_{ij} depends only on S_{ij} , R_{ij} , and δ_{ij} .

The frame-invariant representation in (5.9) could be implemented via autonomic closure, and would thereby allow the coefficients c_α to vary point-to-point in response to changes in the local turbulence state, independent of any assumed model for how the coefficients might depend on the scalar invariants I_0, \dots, I_6 . However, Lund & Novikov (1992) and Doronina *et al* (2018) showed that even when these coefficients are allowed to vary freely, there is little improvement obtained over a constant-coefficient implementation of (5.9). They conclude that the subgrid stress tensor must depend on more than just S_{ij} , R_{ij} , and δ_{ij} and that such additional parametric dependence must be included to substantially improve modeling of the subgrid stress tensor τ_{ij} .

Indeed, in autonomic closure the velocity component values on the $3 \times 3 \times 3$ stencil, relative to the corresponding values at the stencil center point, provide not only the strain and rotation rate tensors $\mathbf{S} \equiv S_{ij}$ and $\mathbf{R} \equiv R_{ij}$, but also their gradients $\nabla \mathbf{S} \equiv \partial S_{ij} / \partial x_k$ and $\nabla \mathbf{R} \equiv \partial R_{ij} / \partial x_k$. In fact, that is *all* the information that is contained in the velocity component values on the $3 \times 3 \times 3$ stencil when translation invariance is enforced by subtracting the stencil-center velocity. Thus, it must be possible to formulate autonomic closure even more generally, and more compactly, as the sum over the complete invariance-preserving set of tensor bases,

analogous to $\mathbf{m}^{(\alpha)}$ in (5.8), that can be formed from \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$, $\nabla\mathbf{R}$, and \mathbf{I} while preserving the rank, symmetry, translation and rotation properties of τ_{ij} . It is such a complete tensor invariant formulation of autonomic closure that is pursued here.

5.3.3 Frame-Invariant Combinations of \mathbf{I} , \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$

This section uses combinations of powers of the following tensors that are available in autonomic closure, namely

$$\mathbf{I} \equiv \delta_{ii} \quad (5.11a)$$

$$\mathbf{S} \equiv S_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \quad (5.11b)$$

$$\mathbf{R} \equiv R_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i} \right) \quad (5.11c)$$

$$\nabla\mathbf{S} = \partial S_{ij} / \partial x_k \quad (5.11d)$$

$$\nabla\mathbf{R} = \partial R_{ij} / \partial x_k \quad (5.11e)$$

which provide the most general complete invariance-preserving tensor basis for representing τ_{ij} under the assumption that the subgrid stress depends only on \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$, $\nabla\mathbf{R}$, and \mathbf{I} .

5.3.3.1 Complete set of rank-2 tensor polynomial bases

The subgrid stress τ_{ij} , which is to be expressed as a polynomial function of the tensors in (5.10), is of rank-2. While \mathbf{S} and \mathbf{R} are rank-2 tensors, $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$ are rank-3 tensors and, due to the lack of an equivalent of the Cayley-Hamilton theorem for rank-3 tensors, there is no general frame-invariant formulation of a tensor polynomial for a rank-2 tensor in terms of a combination of rank-2 and rank-3 tensors. However Smith (1971), Pennisi and Trovato (1987), Zheng (1994), and Itskov (2007) provide a general formulation for expressing any rank-2 tensor \mathbf{B} in the complete and minimal frame-invariant tensor polynomial basis that can be formed from any finite set of rank-2 tensors \mathbf{A}_k , where $k = 1, 2, \dots, N$.

Specifically, each of the N tensors \mathbf{A}_k is first separated into its symmetric and anti-symmetric parts as

$$\text{symmetric part:} \quad \mathbf{M}_i \equiv \frac{1}{2} (\mathbf{A}_k + \mathbf{A}_k^T) \quad i = 1, 2, \dots, m \leq N \quad (5.12a)$$

$$\text{anti-symmetric part: } \mathbf{W}_p \equiv \frac{1}{2}(\mathbf{A}_k - \mathbf{A}_k^T) \quad p = 1, 2, \dots, w \leq N \quad (5.12b)$$

Similarly, $\mathbf{B} = \mathbf{f}(\mathbf{A}_k)$ is separated into symmetric and anti-symmetric parts as

$$\text{symmetric part: } \mathbf{B}_S \equiv \frac{1}{2}(\mathbf{B} + \mathbf{B}^T) \quad (5.13a)$$

$$\text{anti-symmetric part: } \mathbf{B}_A \equiv \frac{1}{2}(\mathbf{B} - \mathbf{B}^T) \quad (5.13b)$$

from which \mathbf{B} can then be reconstructed as $\mathbf{B} = \mathbf{B}_S + \mathbf{B}_A$. Following Smith (1971), Pennisi and Trovato (1987), and Itskov (2007), the symmetric part \mathbf{B}_S can be a function only of the following invariant symmetric rank-2 tensor polynomial bases

$$\mathbf{m}_S^{(0)} = \mathbf{I} \quad (5.14a)$$

$$\mathbf{m}_S^{(1,i)} = \mathbf{M}_i \quad (5.14b)$$

$$\mathbf{m}_S^{(2,i)} = \mathbf{M}_i^2 \quad (5.14c)$$

$$\mathbf{m}_S^{(3,ij)} = \mathbf{M}_i \mathbf{M}_j + \mathbf{M}_j \mathbf{M}_i \quad (5.14d)$$

$$\mathbf{m}_S^{(4,ij)} = \mathbf{M}_i^2 \mathbf{M}_j + \mathbf{M}_j \mathbf{M}_i^2 \quad (5.14e)$$

$$\mathbf{m}_S^{(5,ij)} = \mathbf{M}_i \mathbf{M}_j^2 + \mathbf{M}_j^2 \mathbf{M}_i \quad (5.14f)$$

$$\mathbf{m}_S^{(6,p)} = \mathbf{W}_p^2 \quad (5.14g)$$

$$\mathbf{m}_S^{(7,pq)} = \mathbf{W}_p \mathbf{W}_q - \mathbf{W}_q \mathbf{W}_p \quad (5.14h)$$

$$\mathbf{m}_S^{(8,pq)} = \mathbf{W}_p^2 \mathbf{W}_q - \mathbf{W}_q \mathbf{W}_p^2 \quad (5.14i)$$

$$\mathbf{m}_S^{(9,pq)} = \mathbf{W}_p \mathbf{W}_q^2 - \mathbf{W}_q^2 \mathbf{W}_p \quad (5.14j)$$

$$\mathbf{m}_S^{(10,ip)} = \mathbf{M}_i \mathbf{W}_p - \mathbf{W}_p \mathbf{M}_i \quad (5.14k)$$

$$\mathbf{m}_S^{(11,ip)} = \mathbf{W}_p \mathbf{M}_i \mathbf{W}_p \quad (5.14l)$$

$$\mathbf{m}_S^{(12,ip)} = \mathbf{M}_i^2 \mathbf{W}_p - \mathbf{W}_p \mathbf{M}_i^2 \quad (5.14m)$$

$$\mathbf{m}_S^{(13,ip)} = \mathbf{W}_p \mathbf{M}_i \mathbf{W}_p^2 - \mathbf{W}_p^2 \mathbf{M}_i \mathbf{W}_p \quad (5.14n)$$

for all $i < j = 1, 2, \dots, m$ and all $p < q = 1, 2, \dots, w$. Similarly, the anti-symmetric part \mathbf{B}_A can be a function only of the following invariant rank-2 anti-symmetric tensor polynomial

bases

$$\mathbf{m}_A^{(1,i)} = \mathbf{W}_p \quad (5.15a)$$

$$\mathbf{m}_A^{(2,pq)} = \mathbf{W}_p \mathbf{W}_q - \mathbf{W}_q \mathbf{W}_p \quad (5.15b)$$

$$\mathbf{m}_A^{(3,ij)} = \mathbf{M}_i \mathbf{M}_j - \mathbf{M}_j \mathbf{M}_i \quad (5.15c)$$

$$\mathbf{m}_A^{(4,ij)} = \mathbf{M}_i^2 \mathbf{M}_j - \mathbf{M}_j \mathbf{M}_i^2 \quad (5.15d)$$

$$\mathbf{m}_A^{(5,ij)} = \mathbf{M}_i \mathbf{M}_j^2 - \mathbf{M}_j^2 \mathbf{M}_i \quad (5.15e)$$

$$\mathbf{m}_A^{(6,ij)} = \mathbf{M}_i \mathbf{M}_j \mathbf{M}_i^2 - \mathbf{M}_i^2 \mathbf{M}_j \mathbf{M}_i \quad (5.15f)$$

$$\mathbf{m}_A^{(7,ij)} = \mathbf{M}_j \mathbf{M}_i \mathbf{M}_j^2 - \mathbf{M}_j^2 \mathbf{M}_i \mathbf{M}_j \quad (5.15g)$$

$$\begin{aligned} \mathbf{m}_A^{(8,ijk)} &= \mathbf{M}_i \mathbf{M}_j \mathbf{M}_k + \mathbf{M}_j \mathbf{M}_k \mathbf{M}_i + \mathbf{M}_k \mathbf{M}_i \mathbf{M}_j \\ &\quad - \mathbf{M}_j \mathbf{M}_i \mathbf{M}_k - \mathbf{M}_i \mathbf{M}_k \mathbf{M}_j - \mathbf{M}_k \mathbf{M}_j \mathbf{M}_i \end{aligned} \quad (5.15h)$$

$$\mathbf{m}_A^{(9,ip)} = \mathbf{M}_i \mathbf{W}_p + \mathbf{W}_p \mathbf{M}_i \quad (5.15i)$$

$$\mathbf{m}_A^{(10,ip)} = \mathbf{M}_i \mathbf{W}_p^2 - \mathbf{W}_p^2 \mathbf{M}_i \quad (5.15j)$$

for all $i < j = 1, 2, \dots, m$ and all $p < q = 1, 2, \dots, w$.

Smith (1971) asserts the symmetric and anti-symmetric tensor polynomial bases in (5.14a-n) and (5.15a-j) to be *complete*, meaning that *any* \mathbf{B}_S and \mathbf{B}_A can be written as a linear sum of the corresponding polynomial terms $\mathbf{m}_S^{(\alpha)}$ and $\mathbf{m}_A^{(\alpha)}$, each weighted by a corresponding coefficient. However, a complete tensor polynomial basis is *not minimal* if it is reducible to an even smaller basis set that suffices to represent any rank-2 polynomial \mathbf{B} . This arises from the fact that there may be tensor polynomial relations among various terms in the basis set that allow the number of tensor products in the basis set to be further reduced. Such relations are generally called *Rivlin identities*, and they result from the generalized Cayley-Hamilton theorem

$$\mathbf{A}_k^n - I_A^{(1)} \mathbf{A}_k^{(n-1)} + I_A^{(2)} \mathbf{A}_k^{(n-2)} + \dots + (-1)^n I_A^{(n)} \mathbf{I} = 0 \quad (5.16)$$

where the $I_A^{(i)}$ are scalar invariants of \mathbf{A} defined as $I_A^{(1)} = \text{tr}(\mathbf{A})$, $2I_A^{(2)} = \text{tr}(\mathbf{A})^2 - \text{tr}(\mathbf{A}^2)$, ..., $nI = \det(\mathbf{A})$. By differentiating (5.16) repeatedly with respect to \mathbf{A} , numerous Rivlin

identities can be generated in the form of relations among tensor products of various orders. The set of possible Rivlin identities is very large, making it difficult to prove that a given tensor polynomial basis set is minimal.

Pennisi & Trovato (1987) first proved the irreducibility of Smith's (1971) tensor bases in (5.14) and (5.15), thereby establishing these as a *complete and minimal basis*. Prior to that, a number of rank-2 tensor polynomial bases had been proposed that were complete but were not minimal. Zemach (1998) discusses completeness and minimality of tensor polynomial bases, and Itskov (2007) uses modern tensor notation and algebra to more clearly derive the complete and minimal bases in (5.13) and (5.14). Even complete and minimal bases may not appear unique, since Rivlin identities may allow terms in one basis to be expressed equivalently but differently in another basis. The minimality of a basis simply means that there is no other basis that can be complete and have a smaller number of basis tensors $\mathbf{m}^{(k)}$.

5.3.3.2 The tensor basis in (5.13) and (5.14) is minimal

This section shows that the tensor polynomial basis in (5.8), which has been widely used for representing the turbulent stress tensor τ_{ij} , is not minimal. To do this, the complete and minimal tensor bases in (5.13) and (5.14) are applied to the special case considered in Section 5.3.1.2, for which τ_{ij} is taken to depend only on \mathbf{S} and \mathbf{R} . In that case, in (5.12a,b) $\mathbf{M}_i \equiv \mathbf{M}_1 = \mathbf{S}$ and $\mathbf{W}_p \equiv \mathbf{W}_1 = \mathbf{R}$, and (5.14) then shows that the *complete and minimal* symmetric tensor basis for τ_{ij} as

$$\begin{aligned}
 \mathbf{m}_S^{(0)} &= \mathbf{I} & \mathbf{m}_S^{(10)} &= \mathbf{SR} - \mathbf{RS} \\
 \mathbf{m}_S^{(1)} &= \mathbf{S} & \mathbf{m}_S^{(11)} &= \mathbf{RSR} \\
 \mathbf{m}_S^{(2)} &= \mathbf{S}^2 & \mathbf{m}_S^{(12)} &= \mathbf{S}^2\mathbf{R} - \mathbf{RS}^2 \\
 \mathbf{m}_S^{(6)} &= \mathbf{R}^2 & \mathbf{m}_S^{(13)} &= \mathbf{RSR}^2 - \mathbf{R}^2\mathbf{SR}
 \end{aligned} \tag{5.17}$$

with all other $\mathbf{m}_S^{(k)} \equiv 0$. These should be equivalent to (5.8a-k) if both tensor bases are complete and minimal. However, even then they need not appear identical, since there may be Rivlin identities that allow them to be rearranged into identical forms.

The above $\mathbf{m}_S^{(k)}$ for $k=0, 1, 2, 6, 10, 12,$ and 13 are indeed identical to (5.8a), (5.8b), (5.8c), (5.8d), (5.8e), (5.8f), and (5.8j). However, in (5.8a-k) there is no obvious equivalent to the $\mathbf{m}_S^{(11)}$ term above, nor are there terms above that appear obviously equivalent to (5.8g), (5.8i), or (5.8k). There may be Rivlin identities that allow (5.8g), (5.8i), or (5.8k) to be constructed from the eight symmetric basis tensors $\mathbf{m}_S^{(k)}$ above. However, since the above basis set $\mathbf{m}_S^{(k)}$ involves just eight tensors, while the basis set $\mathbf{m}^{(k)}$ in (5.8a-k) involves ten tensors, the basis set in (5.8a-k) cannot be minimal.

In fluid dynamics, the basis in (5.8a-k) was first introduced by Lumley (1970) and then adopted, with minor corrections, by Pope (1975). Lund & Novikov (1992) and Gatski & Speziale (1993) adopted Pope's (1975) basis. However Pope appears not to have been aware that his basis is not minimal, and to have been unaware of the more general complete basis formulation by Smith (1971) in (5.14a-n) and (5.15a-j). Later, Pennisi & Torvato (1987) proved that Smith's bases are minimal. Even today, it is not widely known in the fluid dynamics community that (5.8) is not a minimal basis, and that by contrast the eight symmetric basis tensors $\mathbf{m}_S^{(k)}$ listed above are a *complete* and *minimal* basis for representing τ_{ij} solely in terms of \mathbf{S} and \mathbf{R} .

5.3.3.3 Tensor elements \mathbf{A}_k in terms of \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$.

Taking $\mathbf{B} \equiv B_{ij}$ as the subgrid stress τ_{ij} , which is symmetric, only the symmetric tensor polynomial basis in (5.14a-n) is needed in a tensor polynomial representation of τ_{ij} . \mathbf{S} and \mathbf{R} are respectively symmetric and anti-symmetric rank-2 tensors, and thus can be included directly in the set of tensors \mathbf{A}_k in Section 5.3.3.1. Thus

$$\mathbf{A}_1 = \mathbf{S} \tag{5.18a}$$

$$\mathbf{A}_2 = \mathbf{R} \tag{5.18b}$$

However, $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$ are rank-3 tensors and, since there does not appear to be a general formulation of a rank-2 tensor polynomial in terms of rank-2 and rank-3 tensors, this study

chooses to form rank-2 tensors from $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$, and include these in the set of tensors in Section II.A. To do this, consider tensor contractions of the form of $\nabla\mathbf{S}^\alpha$, $\nabla\mathbf{R}^\beta$, and $\nabla\mathbf{S}^\gamma\nabla\mathbf{R}^\delta$.

5.3.3.3.1 Rank-2 contractions involving only $\nabla\mathbf{S}^\alpha$

Since $\nabla\mathbf{S}$ is a rank-3 tensor, this section determines the powers α that contract $\nabla\mathbf{S}^\alpha$ to a rank-2 tensor. For any α , $\nabla\mathbf{S}^\alpha$ involves 3α indices, so two of these must be the free indices i and j , and the remaining indices must be repeated in integer m pairs. Thus allowable values of $\alpha > 1$ must satisfy $3\alpha - 2 = 2m$ for integer $m \geq 1$, which is the case only for $\alpha = 2, 4, 6, \dots$. For $\alpha = 2$ there are eight possible tensor products of the form $\nabla\mathbf{S}^2$, namely

$$\begin{aligned} & \frac{\partial S_{kk}}{\partial x_i} \frac{\partial S_{ll}}{\partial x_j}, \quad \frac{\partial S_{kl}}{\partial x_i} \frac{\partial S_{kl}}{\partial x_j}, \quad \frac{\partial S_{ik}}{\partial x_k} \frac{\partial S_{lj}}{\partial x_l}, \quad \frac{\partial S_{ik}}{\partial x_l} \frac{\partial S_{lj}}{\partial x_k} \\ & \frac{\partial S_{ik}}{\partial x_k} \frac{\partial S_{ll}}{\partial x_j}, \quad \frac{\partial S_{ik}}{\partial x_l} \frac{\partial S_{kl}}{\partial x_j}, \quad \frac{\partial S_{kk}}{\partial x_i} \frac{\partial S_{lj}}{\partial x_l}, \quad \frac{\partial S_{kl}}{\partial x_i} \frac{\partial S_{kj}}{\partial x_l} \end{aligned} \quad (5.19)$$

For each of $\alpha = 4, 6, \dots$ there are *far* larger numbers of tensor products of the form $\nabla\mathbf{S}^\alpha$ that can form rank-2 symmetric tensors. Many of these may be reducible via the equivalent of Rivlin identities among them, though for rank-3 tensors there appears to be no equivalent of the Cayley-Hamilton theorem from which to obtain such identities. Such an approach could potentially lead to a *minimal* tensor polynomial basis set. However, while there are efficiencies gained from a tensor representation in a minimal basis, there is no loss of generality if a non-minimal basis is used. Moreover, lacking the equivalent of a Cayley-Hamilton theorem for rank-3 tensors leaves open the question of whether it is even possible for there to be a *finite* set of tensor products $\nabla\mathbf{S}^\alpha$ that can form rank-2 symmetric tensors.

Thus, while including these higher-power tensor products $\nabla\mathbf{S}^\alpha$ for $\alpha > 2$ may be necessary to obtain a *complete* tensor polynomial basis set, even if only the $\nabla\mathbf{S}^2$ products are included in the representation of τ_{ij} they will allow effects of $\nabla\mathbf{S}$ to be reflected in the subgrid stress – though in less than a fully complete way. The complete set of unique second-order rank-2 contractions $\nabla\mathbf{S}^2$ consists of

$$\mathbf{A}_3 = \frac{\partial S_{kk}}{\partial x_i} \frac{\partial S_{ll}}{\partial x_j} \quad (5.20a)$$

$$\mathbf{A}_4 = \frac{\partial S_{kl}}{\partial x_i} \frac{\partial S_{kl}}{\partial x_j} \quad (5.20b)$$

$$\mathbf{A}_5 = \frac{\partial S_{ik}}{\partial x_k} \frac{\partial S_{lj}}{\partial x_l} \quad (5.20c)$$

$$\mathbf{A}_6 = \frac{\partial S_{ik}}{\partial x_l} \frac{\partial S_{lj}}{\partial x_k} \quad (5.20d)$$

$$\mathbf{A}_7 = \frac{\partial S_{ik}}{\partial x_k} \frac{\partial S_{lj}}{\partial x_j} \quad (5.20e)$$

$$\mathbf{A}_8 = \frac{\partial S_{ik}}{\partial x_l} \frac{\partial S_{kl}}{\partial x_j} \quad (5.20f)$$

$$\mathbf{A}_9 = \frac{\partial S_{kk}}{\partial x_i} \frac{\partial S_{lj}}{\partial x_l} \quad (5.20g)$$

$$\mathbf{A}_{10} = \frac{\partial S_{kl}}{\partial x_i} \frac{\partial S_{kj}}{\partial x_l} \quad (5.20h)$$

5.3.3.3.2 Rank-2 contractions involving only $\nabla \mathbf{R}^\beta$

Because in Section 5.3.3.1 the tensors \mathbf{A}_k are identified without regard to their symmetry, since $\nabla \mathbf{R}$ is a rank-3 tensor (like $\nabla \mathbf{S}$) it is possible by direct analogy with the rank-2 contractions $\nabla \mathbf{S}^2$ in (5.20a-h) to write corresponding rank-2 contractions $\nabla \mathbf{R}^2$. However, due to the anti-symmetry of \mathbf{R} it is always the case that $R_{kk} \equiv 0$, even when $S_{kk} \neq 0$, and this eliminates three of the eight corresponding contractions $\nabla \mathbf{R}^2$, leaving only

$$\mathbf{A}_{11} = \frac{\partial R_{kl}}{\partial x_i} \frac{\partial R_{kl}}{\partial x_j} \quad (5.21a)$$

$$\mathbf{A}_{12} = \frac{\partial R_{ik}}{\partial x_k} \frac{\partial R_{lj}}{\partial x_l} \quad (5.21b)$$

$$\mathbf{A}_{13} = \frac{\partial R_{ik}}{\partial x_l} \frac{\partial R_{lj}}{\partial x_k} \quad (5.21c)$$

$$\mathbf{A}_{14} = \frac{\partial R_{ik}}{\partial x_l} \frac{\partial R_{kl}}{\partial x_j} \quad (5.21d)$$

$$\mathbf{A}_{15} = \frac{\partial R_{kl}}{\partial x_i} \frac{\partial R_{kj}}{\partial x_l} \quad (5.21e)$$

5.3.3.3.3 Rank-2 contractions involving $\nabla \mathbf{S}^\gamma \nabla \mathbf{R}^\delta$

Similarly, because $\nabla \mathbf{S}$ and $\nabla \mathbf{R}$ are both rank-3 tensors it is again possible by direct analogy with the rank-2 contractions in (5.20a-h) to write corresponding second-order rank-2 mixed contractions $\nabla \mathbf{S} \nabla \mathbf{R}$. As above, contractions involving R_{kk} are eliminated since in all cases

$R_{kk} \equiv 0$, but in this case this eliminates only two of the eight contractions. Additionally, the symmetry of \mathbf{S} and the anti-symmetry of \mathbf{R} require $S_{kl}R_{kl} \equiv 0$, but due to resulting chain-rule terms this does not eliminate any of the remaining six contractions. Thus, retaining only second-order rank-2 mixed contractions $\nabla\mathbf{S}\nabla\mathbf{R}$ leaves

$$\mathbf{A}_{16} = \frac{\partial S_{kl}}{\partial x_i} \frac{\partial R_{kl}}{\partial x_j} \quad (5.22a)$$

$$\mathbf{A}_{17} = \frac{\partial S_{ik}}{\partial x_k} \frac{\partial R_{lj}}{\partial x_l} \quad (5.22b)$$

$$\mathbf{A}_{18} = \frac{\partial S_{ik}}{\partial x_l} \frac{\partial R_{lj}}{\partial x_k} \quad (5.22c)$$

$$\mathbf{A}_{19} = \frac{\partial S_{ik}}{\partial x_l} \frac{\partial R_{kl}}{\partial x_j} \quad (5.22d)$$

$$\mathbf{A}_{20} = \frac{\partial S_{kk}}{\partial x_i} \frac{\partial R_{lj}}{\partial x_l} \quad (5.22e)$$

$$\mathbf{A}_{21} = \frac{\partial S_{kl}}{\partial x_i} \frac{\partial R_{kj}}{\partial x_l} \quad (5.22f)$$

5.3.3.4 Symmetric and antisymmetric tensors \mathbf{M}_k and \mathbf{W}_k

Next, the symmetric and anti-symmetric parts of each of the tensor elements \mathbf{A}_k in (5.18), (5.20), (5.21), and (5.22) are formed via (5.12a,b). Some of the resulting symmetric parts \mathbf{M}_k are duplicates upon addition in (5.12a), and thus are listed only once, and some of the resulting anti-symmetric parts \mathbf{W}_k are zero upon subtraction in (5.12b). The resulting unique symmetric parts are

$$\mathbf{M}_1 = S_{ij} \quad (5.23a)$$

$$\mathbf{M}_2 = \frac{\partial S_{kk}}{\partial x_i} \frac{\partial S_{ll}}{\partial x_j} \quad (5.23b)$$

$$\mathbf{M}_3 = \frac{\partial S_{kl}}{\partial x_i} \frac{\partial S_{kl}}{\partial x_j} \quad (5.23c)$$

$$\mathbf{M}_4 = \frac{\partial S_{ik}}{\partial x_k} \frac{\partial S_{lj}}{\partial x_l} \quad (5.23d)$$

$$\mathbf{M}_5 = \frac{\partial S_{ik}}{\partial x_l} \frac{\partial S_{lj}}{\partial x_k} \quad (5.23e)$$

$$\mathbf{M}_6 = \frac{1}{2} \left(\frac{\partial S_{ik}}{\partial x_k} \frac{\partial S_{ll}}{\partial x_j} + \frac{\partial S_{ll}}{\partial x_j} \frac{\partial S_{ik}}{\partial x_k} \right) \quad (5.23f)$$

$$\mathbf{M}_7 = \frac{1}{2} \left(\frac{\partial S_{ik}}{\partial x_l} \frac{\partial S_{kl}}{\partial x_j} + \frac{\partial S_{kl}}{\partial x_i} \frac{\partial S_{jk}}{\partial x_l} \right) \quad (5.23g)$$

$$\mathbf{M}_8 = \frac{\partial R_{kl}}{\partial x_i} \frac{\partial R_{kl}}{\partial x_j} \quad (5.23h)$$

$$\mathbf{M}_9 = \frac{\partial R_{ik}}{\partial x_k} \frac{\partial R_{lj}}{\partial x_l} \quad (5.23i)$$

$$\mathbf{M}_{10} = \frac{\partial R_{ik}}{\partial x_k} \frac{\partial R_{lj}}{\partial x_l} \quad (5.23j)$$

$$\mathbf{M}_{11} = \frac{1}{2} \left(\frac{\partial R_{ik}}{\partial x_l} \frac{\partial R_{kl}}{\partial x_j} + \frac{\partial R_{kl}}{\partial x_i} \frac{\partial R_{jk}}{\partial x_l} \right) \quad (5.23k)$$

$$\mathbf{M}_{12} = \frac{1}{2} \left(\frac{\partial S_{ik}}{\partial x_k} \frac{\partial R_{lj}}{\partial x_l} + \frac{\partial R_{li}}{\partial x_l} \frac{\partial S_{jk}}{\partial x_k} \right) \quad (5.23l)$$

$$\mathbf{M}_{13} = \frac{1}{2} \left(\frac{\partial S_{ik}}{\partial x_l} \frac{\partial R_{lj}}{\partial x_k} + \frac{\partial R_{li}}{\partial x_k} \frac{\partial S_{jk}}{\partial x_l} \right) \quad (5.23m)$$

$$\mathbf{M}_{14} = \frac{1}{2} \left(\frac{\partial S_{kk}}{\partial x_i} \frac{\partial R_{lj}}{\partial x_l} + \frac{\partial R_{li}}{\partial x_l} \frac{\partial S_{kk}}{\partial x_j} \right) \quad (5.23n)$$

$$\mathbf{M}_{15} = \frac{1}{2} \left(\frac{\partial S_{kl}}{\partial x_i} \frac{\partial R_{kj}}{\partial x_l} + \frac{\partial R_{ki}}{\partial x_l} \frac{\partial S_{kl}}{\partial x_j} \right) \quad (5.23o)$$

Thus there are $m = 15$ symmetric parts that can be formed from the tensor elements \mathbf{A}_k in (5.19)-(5.21). Similarly, the resulting unique non-zero anti-symmetric parts are

$$\mathbf{W}_1 = R_{ij} \quad (5.24a)$$

$$\mathbf{W}_2 = \frac{1}{2} \left(\frac{\partial S_{ik}}{\partial x_k} \frac{\partial S_{ll}}{\partial x_j} - \frac{\partial S_{ll}}{\partial x_i} \frac{\partial S_{jk}}{\partial x_k} \right) \quad (5.24b)$$

$$\mathbf{W}_3 = \frac{1}{2} \left(\frac{\partial S_{ik}}{\partial x_l} \frac{\partial S_{kl}}{\partial x_j} - \frac{\partial S_{kl}}{\partial x_i} \frac{\partial S_{jk}}{\partial x_l} \right) \quad (5.24c)$$

$$\mathbf{W}_4 = \frac{1}{2} \left(\frac{\partial S_{kk}}{\partial x_i} \frac{\partial S_{lj}}{\partial x_l} - \frac{\partial S_{li}}{\partial x_l} \frac{\partial S_{kk}}{\partial x_j} \right) \quad (5.24d)$$

$$\mathbf{W}_5 = \frac{1}{2} \left(\frac{\partial R_{ik}}{\partial x_l} \frac{\partial R_{kl}}{\partial x_j} - \frac{\partial R_{kl}}{\partial x_i} \frac{\partial R_{jk}}{\partial x_l} \right) \quad (5.24e)$$

$$\mathbf{W}_6 = \frac{1}{2} \left(\frac{\partial R_{kl}}{\partial x_i} \frac{\partial R_{kj}}{\partial x_l} - \frac{\partial R_{ki}}{\partial x_l} \frac{\partial R_{kl}}{\partial x_j} \right) \quad (5.24f)$$

$$\mathbf{W}_7 = \frac{1}{2} \left(\frac{\partial S_{kl}}{\partial x_i} \frac{\partial R_{kl}}{\partial x_j} - \frac{\partial R_{kl}}{\partial x_i} \frac{\partial S_{kl}}{\partial x_j} \right) \quad (5.24g)$$

$$\mathbf{W}_8 = \frac{1}{2} \left(\frac{\partial S_{ik}}{\partial x_k} \frac{\partial R_{lj}}{\partial x_l} - \frac{\partial R_{li}}{\partial x_l} \frac{\partial S_{jk}}{\partial x_k} \right) \quad (5.24h)$$

$$\mathbf{W}_9 = \frac{1}{2} \left(\frac{\partial S_{ik}}{\partial x_l} \frac{\partial R_{lj}}{\partial x_k} - \frac{\partial R_{li}}{\partial x_k} \frac{\partial S_{jk}}{\partial x_l} \right) \quad (5.24i)$$

$$\mathbf{W}_{10} = \frac{1}{2} \left(\frac{\partial S_{ik}}{\partial x_l} \frac{\partial R_{kl}}{\partial x_j} - \frac{\partial R_{kl}}{\partial x_i} \frac{\partial S_{jk}}{\partial x_l} \right) \quad (5.24j)$$

$$\mathbf{W}_{11} = \frac{1}{2} \left(\frac{\partial S_{kk}}{\partial x_i} \frac{\partial R_{lj}}{\partial x_l} - \frac{\partial R_{li}}{\partial x_l} \frac{\partial S_{kk}}{\partial x_j} \right) \quad (5.24k)$$

$$\mathbf{W}_{12} = \frac{1}{2} \left(\frac{\partial S_{kl}}{\partial x_i} \frac{\partial R_{kj}}{\partial x_l} - \frac{\partial R_{ki}}{\partial x_l} \frac{\partial S_{kl}}{\partial x_j} \right) \quad (5.24l)$$

Thus there are $w = 12$ anti-symmetric parts that can be formed from the tensor elements \mathbf{A}_k in (5.19)-(5.21).

5.3.3.5 Symmetric rank-2 tensor polynomial basis $\mathbf{m}_S^{(i)}$

Following Section 5.3.3.1, any symmetric rank-2 tensor can be expressed as a polynomial in the invariant symmetric rank-2 tensor polynomial basis $\mathbf{m}_S^{(k)}$ given in (5.14a-n). These are based on the second-order symmetric tensors \mathbf{M}_i for $i = 1, 2, \dots, 15$ in (5.22), and on the second-order anti-symmetric tensors \mathbf{W}_p for $p = 1, 2, \dots, 12$ in (5.23). The resulting number of $\mathbf{m}_S^{(k)}$ is large. For instance, the first few $\mathbf{m}_S^{(k,\alpha)}$ for $k = 0, 1, 2, 3, \dots, 13$ are given by

$$k = 0 \quad \mathbf{m}_S^{(0)} = \mathbf{I} \quad (5.25a)$$

$$k = 1 \quad \mathbf{m}_S^{(1,1)} = \mathbf{M}_1 \quad (5.25b)$$

$$\mathbf{m}_S^{(1,2)} = \mathbf{M}_2 \quad (5.25c)$$

⋮

$$\mathbf{m}_S^{(1,15)} = \mathbf{M}_{15} \quad (5.25d)$$

$$k = 2 \quad \mathbf{m}_S^{(2,1)} = \mathbf{M}_1^2 \quad (5.25e)$$

$$\mathbf{m}_S^{(2,2)} = \mathbf{M}_2^2 \quad (5.25f)$$

⋮

$$\mathbf{m}_S^{(2,15)} = \mathbf{M}_{15}^2 \quad (5.25g)$$

$$k = 3 \quad \mathbf{m}_S^{(3,1)} = \mathbf{M}_1 \mathbf{M}_2 + \mathbf{M}_2 \mathbf{M}_1 \quad (5.25h)$$

$$\mathbf{m}_S^{(3,2)} = \mathbf{M}_1 \mathbf{M}_3 + \mathbf{M}_3 \mathbf{M}_1 \quad (5.25i)$$

$$\mathbf{m}_S^{(3,3)} = \mathbf{M}_1 \mathbf{M}_4 + \mathbf{M}_4 \mathbf{M}_1 \quad (5.25j)$$

$$\begin{aligned} & \vdots \\ \mathbf{m}_S^{(3,14)} &= \mathbf{M}_1\mathbf{M}_{15} + \mathbf{M}_{15}\mathbf{M}_1 \end{aligned} \quad (5.25k)$$

$$\mathbf{m}_S^{(3,15)} = \mathbf{M}_2\mathbf{M}_3 + \mathbf{M}_3\mathbf{M}_2 \quad (5.25l)$$

$$\mathbf{m}_S^{(3,16)} = \mathbf{M}_2\mathbf{M}_4 + \mathbf{M}_4\mathbf{M}_2 \quad (5.25m)$$

$$\mathbf{m}_S^{(3,17)} = \mathbf{M}_2\mathbf{M}_5 + \mathbf{M}_5\mathbf{M}_2 \quad (5.25n)$$

$$\begin{aligned} & \vdots \\ \mathbf{m}_S^{(3,30)} &= \mathbf{M}_2\mathbf{M}_{15} + \mathbf{M}_{15}\mathbf{M}_2 \end{aligned} \quad (5.25o)$$

$$\begin{aligned} & \vdots \\ \mathbf{m}_S^{(3,108)} &= \mathbf{M}_{14}\mathbf{M}_{15} + \mathbf{M}_{15}\mathbf{M}_{14} \end{aligned} \quad (5.25p)$$

It is easy to determine how many $\mathbf{m}_S^{(k,\alpha)}$ there are, given that $i < j = 1, 2, \dots, 15$ and $p < q = 1, 2, \dots, 12$. Specifically

$$\mathbf{m}_S^{(0)} : 1 \quad (5.26a)$$

$$\mathbf{m}_S^{(1,\alpha)} : 15 \quad (5.26b)$$

$$\mathbf{m}_S^{(2,\alpha)} : 15 \quad (5.26c)$$

$$\mathbf{m}_S^{(3,\alpha)} : (15^2 - 15)/2 = 105 \quad (5.26d)$$

$$\mathbf{m}_S^{(4,\alpha)} : (15^2 - 15)/2 = 105 \quad (5.26e)$$

$$\mathbf{m}_S^{(5,\alpha)} : (15^2 - 15)/2 = 105 \quad (5.26f)$$

$$\mathbf{m}_S^{(6,\alpha)} : 12 \quad (5.26g)$$

$$\mathbf{m}_S^{(7,\alpha)} : (12^2 - 12)/2 = 66 \quad (5.26h)$$

$$\mathbf{m}_S^{(8,\alpha)} : (12^2 - 12)/2 = 66 \quad (5.26i)$$

$$\mathbf{m}_S^{(9,\alpha)} : (12^2 - 12)/2 = 66 \quad (5.26j)$$

$$\mathbf{m}_S^{(10,\alpha)} : 15 \cdot 12 = 180 \quad (5.26k)$$

$$\mathbf{m}_S^{(11,\alpha)} : 15 \cdot 12 = 180 \quad (5.26l)$$

$$\mathbf{m}_S^{(12,\alpha)} : 15 \cdot 12 = 180 \quad (5.26m)$$

$$\mathbf{m}_S^{(13,\alpha)} : 15 \cdot 12 = 180 \quad (5.26n)$$

Thus the total number of invariant symmetric rank-2 tensor polynomial bases $\mathbf{m}_S^{(k)}$ in (5.14a-n) is 1276, and a complete frame-invariant representation for τ_{ij} in autonomic closure would therefore involve determining a total of 1276 coefficients.

In an “on-the-fly” implementation of autonomic closure these coefficients would need to be determined at each point and time, which is likely to be unacceptably burdensome from a computational perspective. Alternatively, in a “static” implementation of autonomic closure these 1570 coefficients need to be determined just once, in advance of the simulation itself, and can then be utilized without excessive computational burden in a simulation. In both cases, however, the larger computational burden may come from the large number of tensor component multiplications required to calculate the $\mathbf{m}_S^{(k)}$ in (5.14a-n) via the \mathbf{M}_i in (5.23a-r) and the \mathbf{W}_p (5.24a-l).

5.3.3.6 Incompressible case: Deviatoric basis $\mathbf{m}_S^{(i)}$

For incompressible flow, which can be formulated in the deviatoric stress and deviatoric strain rate, many of the $\mathbf{m}_S^{(k)}$ in (5.13) and (5.25) are zero, allowing considerable reduction in the total number of coefficients. In that case $S_{ii} = 0$, so it is possible to work in the deviatoric strain rate \mathbf{S}^{dev} , for which $S_{ii}^{dev} \equiv 0$, and in the deviatoric stress τ^{dev} , for which $\tau^{dev} \equiv 0$ and thus $\mathbf{m}_S^{(0)}$ in (5.14a) and (5.26a) does not appear, as was seen in going from (5.5) to (5.9). The deviatoric stress τ^{dev} can then be formulated in terms of \mathbf{S}^{dev} , \mathbf{R} , $\nabla \mathbf{S}^{dev}$, and $\nabla \mathbf{R}$, which is equivalent to using the results above but enforcing $S_{ii} \equiv 0$ and dropping $\mathbf{m}_S^{(0)}$. As a result, several of the \mathbf{M}_i and \mathbf{W}_p in (5.22) and (5.23) are zero, which reduces the number of polynomial bases $\mathbf{m}_S^{(k)}$.

Specifically, \mathbf{M}_2 in (5.23b), \mathbf{M}_6 in (5.23f), and \mathbf{M}_{14} in (5.23q) are all zero, as are \mathbf{W}_2 in (5.24b), \mathbf{W}_4 in (5.24d), and \mathbf{W}_{11} in (5.24k). This leaves a total of 12 non-zero \mathbf{M}_i and 9 non-zero \mathbf{W}_p . Then, with $i < j = 1, 2, \dots, 12$ and all $p < q = 1, 2, \dots, 9$, following the same

procedure as in (5.26a-n) the total number of invariant symmetric rank-2 tensor polynomial bases $\mathbf{m}_S^{(k)}$ in (5.14a-n) for the incompressible case is 784.

5.3.4 Second-Order Truncation of the Complete Tensor Basis

The invariant symmetric rank-2 tensor polynomial basis $\mathbf{m}_S^{(\alpha)}$ for $\alpha = 0, 1, \dots, 1275$ in (5.14) is *complete* under the assumption that τ_{ij} depends on \mathbf{I} , \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}^2$, $\nabla\mathbf{R}^2$, and $\nabla\mathbf{S}\nabla\mathbf{R}$. Thus all representations of τ_{ij} that satisfy the invariance properties of a symmetric rank-2 tensor can be written as a tensor polynomial

$$\tau_{ij} = \sum_{\alpha=0}^{1275} c_{\alpha} \mathbf{m}_S^{(\alpha)}. \quad (5.27)$$

In (5.27), from (5.13) with (5.22) and (5.23) note that

- $\mathbf{m}_S^{(1,i)}$ contains terms up to first order in \mathbf{M}_i and thereby up to *second* order in the velocity components u_i .
- $\mathbf{m}_S^{(2,i)}$, $\mathbf{m}_S^{(3,ij)}$, $\mathbf{m}_S^{(6,p)}$, $\mathbf{m}_S^{(7,pq)}$, and $\mathbf{m}_S^{(10,ip)}$ contain second-order products of \mathbf{M}_i and \mathbf{W}_p , and thus terms up to *fourth* order in the velocity components u_i .
- $\mathbf{m}_S^{(4,ij)}$, $\mathbf{m}_S^{(5,ij)}$, $\mathbf{m}_S^{(8,pq)}$, $\mathbf{m}_S^{(9,ip)}$, $\mathbf{m}_S^{(11,ip)}$ and $\mathbf{m}_S^{(12,ip)}$ contain third-order products of \mathbf{M}_i and \mathbf{W}_p , and thus terms up to *sixth* order in the velocity components u_i .
- $\mathbf{m}_S^{(13,ip)}$ contains fourth-order tensor products of \mathbf{M}_i and \mathbf{W}_p , and thus terms up to *eighth* order in the velocity components u_i .
- There are *no* terms in this complete and minimal tensor representation of τ_{ij} involving terms above *eighth* order in the velocity components u_i .

Chapter 4 shows that truncating a series representation to only retain velocity products up to second order is sufficient to obtain excellent representation of the subgrid stress and the associated subgrid production, regardless whether the second-order products include non-colocated or only colocated stencil points. Anticipating that at least comparable accuracy will be obtained if (5.27) is similarly truncated to retain at most second-order products of velocities, it is necessary to determine where velocity products of various orders appear in each of the tensor bases $\mathbf{m}_S^{(\alpha)}$.

The strain rate \mathbf{S} in (5.11b) and rotation rate \mathbf{R} in (5.11c) are each linear in the velocity components u_i , and therefore \mathbf{M}_1 in (5.23a) and \mathbf{W}_1 in (5.24a) are each linear in the velocities. The remaining \mathbf{M}_i in (5.23b-r) and \mathbf{W}_p in (5.24b-l) all involve rank-2 contractions of the forms $\nabla\mathbf{S}^2$, $\nabla\mathbf{R}^2$, and $\nabla\mathbf{S}\nabla\mathbf{R}$, and therefore are second-order in the velocity components u_i . From these, the order of the velocity products in each of the tensor bases $\mathbf{m}_S^{(\alpha)}$ in (5.14a-n) can be readily determined.

- Zeroth-order velocity products appear only in $\mathbf{m}_S^{(0)} \equiv \mathbf{I}$.
- First-order velocity products appear only in $\mathbf{m}_S^{(1)}$, since only \mathbf{M}_1 and \mathbf{W}_1 are linear in the velocities, and \mathbf{M}_1 appears linearly in $\mathbf{m}_S^{(1)}$ while \mathbf{W}_1 enters only via tensor products of second order or higher in (5.14g-n).
- Second-order velocity products are present only in $\mathbf{m}_S^{(1)}$ for $i > 1$, in $\mathbf{m}_S^{(2)}$ for $i = 1$, in $\mathbf{m}_S^{(6)}$ for $p = 1$, and in $\mathbf{m}_S^{(10)}$ for $i = p = 1$.
- All other terms in (5.26) are of order three or higher in the velocity components.

Thus, the tensor representation in (5.27) can be truncated to retain all terms that are up to second-order in the velocities, while preserving frame invariance, as

$$\tau_{ij} = c_0\mathbf{I} + \underbrace{c_1\mathbf{m}_S^{(1,1)}}_{\mathbf{S}} + \underbrace{\sum_{\alpha=2}^{15} c_{1,\alpha}\mathbf{m}_S^{(1,\alpha)}}_{\nabla\mathbf{S}^2, \nabla\mathbf{R}^2, \nabla\mathbf{S}\nabla\mathbf{R}} + \underbrace{c_{16}\mathbf{m}_S^{(2,1)}}_{\mathbf{S}^2} + \underbrace{c_{17}\mathbf{m}_S^{(6,1)}}_{\mathbf{R}^2} + \underbrace{c_{18}\mathbf{m}_S^{(10,1)}}_{(\mathbf{SR} - \mathbf{RS})} \quad (5.28)$$

The four non-gradient terms in (5.28) can be expressed directly in \mathbf{S} and \mathbf{R} as indicated above, and from (5.14b) $\mathbf{m}_S^{(1,\alpha)}$ in the gradient terms is simply equal to $\mathbf{M}_{(\alpha)}$ in (5.23b-r), allowing (5.28) to be written as

$$\begin{aligned} \tau_{ij} = & c_0\mathbf{I} + c_1\mathbf{S} + c_{16}\mathbf{S}^2 + c_{17}\mathbf{R}^2 + c_{18}(\mathbf{SR} - \mathbf{RS}) \\ & + \underbrace{\sum_{\alpha=2}^7 c_{\alpha}\mathbf{M}_{(\alpha)}}_{\nabla\mathbf{S}^2} + \underbrace{\sum_{\alpha=8}^{11} c_{\alpha}\mathbf{M}_{(\alpha)}}_{\nabla\mathbf{R}^2} + \underbrace{\sum_{\alpha=12}^{15} c_{\alpha}\mathbf{M}_{(\alpha)}}_{\nabla\mathbf{S}\nabla\mathbf{R}} \end{aligned} \quad (5.29)$$

Equation (5.29) with $\mathbf{M}_{(\alpha)}$ in (5.23b-r) is the most general frame-invariant tensor representation for τ_{ij} (under the assumption that the stress depends only on \mathbf{I} , \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}^2$, $\nabla\mathbf{R}^2$, and

$\nabla\mathbf{S}\nabla\mathbf{R}$) that is complete up to second order in the velocity components u_i . It involves only 19 coefficients.

These 19 terms include all multi-point second-order velocity products, which in a velocity-only series representation required 3403 terms. This complete frame-invariant tensor formulation of autonomic closure uses all of the information available in the velocities on a $3 \times 3 \times 3$ stencil, and does so with the smallest possible number of terms, while ensuring frame invariance. The reduction in the number of terms greatly reduces the computational burden of solving the over-determined linear system for the unknown coefficients, however this comes at the expense of having to compute a substantial number of matrix products $\mathbf{M}_{(\alpha)}$ via (5.23b-r).

A further advantage of (5.29), beyond the reduction in the number of coefficients that must be obtained via the autonomic closure (either in an “on-the-fly” implementation or in a “static” implementation), is the fact that this representation is the only possible “tensorally correct and complete” representation up to second order in the velocities. As a result, unlike the *ad hoc* series representation in Chapters 2-4, the tensorally correct and complete representation in (5.27) and (5.29) should have *universal* coefficients if τ_{ij} can be represented by \mathbf{I} , \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}^2$, $\nabla\mathbf{R}^2$, and $\nabla\mathbf{S}\nabla\mathbf{R}$. This could allow a 1276-term static implementation, or even a 19-term static implementation, to be even more accurate than a dynamic implementation of the earlier series representation, and may even allow a static implementation to be as accurate as a far more computationally burdensome “on-the-fly” implementation. The frame-invariant representation in (5.29) may even be sufficiently computationally efficient to allow fully dynamic “on-the-fly” implementation of autonomic closure.

It is readily apparent that truncations of (5.27) analogous to (5.29) that retain terms of higher order in a similarly “tensorally correct and complete” form can be obtained via the same procedure used in this section.

PERFORMANCE OF INVARIANT REPRESENTATIONS

This chapter applies the metrics developed in Chapter 3 to assess the accuracy and efficiency of frame-invariant generalized representations for F_{ij} in autonomic closure, of the type developed in Chapter 5. These include nonparametric representations expressed directly in the primitive variables of the simulation, such as the stencil velocities \mathbf{u}_m , and parametric representations expressed in tensor quantities derived from the stencil velocities such as \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$, and $\nabla\mathbf{R}$. It then uses those results to determine which frame-invariant representation in autonomic closure can provide the greatest accuracy and efficiency in representing subgrid terms such as the subgrid stress $\tau_{ij}(\mathbf{x}, t)$ and the associated subgrid production $P(\mathbf{x}, t)$.

6.1 Frame-Invariant Representations for F_{ij}

This chapter evaluates the following five frame-invariant generalized representations for F_{ij} obtained from the results in Chapter 4, specifically:

- **Case 10:** This is the full velocity-based representation in Equation (5.1) obtained from application of the Smith (1971) general representation theory to express τ_{ij} in terms of the complete set of $m = 1, 2, \dots, P = 27$ stencil velocities \mathbf{u}_m . Thus

$$\tau = c_0\mathbf{I} + \sum_{m=1}^P c_m(\mathbf{u}_m \otimes \mathbf{u}_m) + \sum_{m=1}^P \sum_{n=m+1}^P c_{mn}(\mathbf{u}_m \otimes \mathbf{u}_n + \mathbf{u}_n \otimes \mathbf{u}_m). \quad (6.1)$$

Since it retains the colocated velocity products in the first sum and the non-colocated velocity products in the second sum, it involves $N = 1 + 27 + (27^2 - 27)/2 = 379$ coefficients. Galilean-invariance is enforced by subtracting the stencil center-point velocity from all the stencil velocities, which in principle allows the number of independent coefficients to be reduced to 352, however here the full set of 379 coefficients is retained. (This is herein also referred to as the ‘‘Smith-UN’’ form.)

- **Case 11:** This is the same as Case 10, but omits the non-colocated velocity products in the second sum in Equation (6.1), justified by the observation in Chapter 4 that including non-colocated velocity products provided no apparent increase in accuracy. Thus

$$\boldsymbol{\tau} = c_0 \mathbf{I} + \sum_{m=1}^P c_m (\mathbf{u}_m \otimes \mathbf{u}_m). \quad (6.2)$$

It therefore involves $N = 1 + 26 = 27$ coefficients. (This is herein also referred to as the ‘‘Smith-UC’’ form.)

- **Case 20:** This is the full invariant representation in Equation (5.28), from application of the Smith (1971) general representation theory to express τ_{ij} in terms of \mathbf{S} , \mathbf{R} , $\nabla \mathbf{S}$, and $\nabla \mathbf{R}$, with truncation of all terms involving velocity products of orders higher than two. Thus, as in Equation (5.29),

$$\begin{aligned} \tau_{ij} = & c_0 \mathbf{I} + c_1 \mathbf{S} + c_{16} \mathbf{S}^2 + c_{17} \mathbf{R}^2 + c_{18} (\mathbf{SR} - \mathbf{RS}) \\ & + \underbrace{\sum_{\alpha=2}^7 c_\alpha \mathbf{M}_{(\alpha)}}_{\nabla \mathbf{S}^2} + \underbrace{\sum_{\alpha=8}^{11} c_\alpha \mathbf{M}_{(\alpha)}}_{\nabla \mathbf{R}^2} + \underbrace{\sum_{\alpha=12}^{15} c_\alpha \mathbf{M}_{(\alpha)}}_{\nabla \mathbf{S} \nabla \mathbf{R}} \end{aligned} \quad (6.3)$$

This representation involves 19 coefficients. (It is herein also referred to as the ‘‘Smith-19’’ form.)

- **Case 21:** This is the frame-invariant representation in Equation (5.9) and Equation (5.8), originally proposed by Lumley (1970) and Pope (1974), based on the assumption that τ_{ij} can be expressed solely in terms of \mathbf{S} and \mathbf{R} . Thus

$$\begin{aligned} \tau_{ij} = & c_0 \mathbf{I} + c_1 \mathbf{S} + c_2 \mathbf{S}^2 + c_3 \mathbf{R}^2 + c_4 (\mathbf{SR} - \mathbf{RS}) \\ & + c_5 (\mathbf{S}^2 \mathbf{R} - \mathbf{RS}^2) + c_6 (\mathbf{SR}^2 - \mathbf{R}^2 \mathbf{S}) + c_7 (\mathbf{S}^2 \mathbf{R}^2 - \mathbf{R}^2 \mathbf{S}^2) \\ & + c_8 (\mathbf{SRS}^2 - \mathbf{S}^2 \mathbf{RS}) + c_9 (\mathbf{RSR}^2 - \mathbf{R}^2 \mathbf{SR}) + c_{10} (\mathbf{RS}^2 \mathbf{R}^2 - \mathbf{R}^2 \mathbf{S}^2 \mathbf{R}) \end{aligned} \quad (6.4)$$

As noted in Section 5.3.3.2, this representation is complete but not minimal. It involves 11 coefficients. (This is herein also referred to as the ‘‘Pope-11’’ form.)

- **Case 22:** This is the frame-invariant representation in (5.1), from application of the Smith (1971) general representation theory on the assumption that τ_{ij} can be expressed solely in terms of \mathbf{S} and \mathbf{R} . Thus

$$\begin{aligned} \tau_{ij} = & c_0\mathbf{I} + c_1\mathbf{S} + c_2\mathbf{S}^2 + c_3\mathbf{R}^2 + c_4(\mathbf{SR} - \mathbf{RS}) \\ & + c_5(\mathbf{RSR}) + c_6(\mathbf{S}^2\mathbf{R} - \mathbf{RS}^2) + c_7(\mathbf{RSR}^2 - \mathbf{R}^2\mathbf{SR}) \end{aligned} \quad (6.5)$$

Since this was obtained in Section 5.3.3 from the Smith (1971) general representation theory, unlike Case 21 it is both complete and minimal. It involves only 8 coefficients. (This is herein also referred to as the “ASU-8” form.)

- **Case 23:** This is the same as Case 22, but retains only products up to second-order in the velocities, from the observation in Chapter 4 that including non-colocated velocity products provided no apparent increase in accuracy. Thus

$$\tau_{ij} = c_0\mathbf{I} + c_1\mathbf{S} + c_2\mathbf{S}^2 + c_3\mathbf{R}^2 + c_4(\mathbf{SR} - \mathbf{RS}) \quad (6.6)$$

This representation involves five coefficients. (It is herein also referred to as the “Smith-5” form.)

6.2 Error-Free Differentiation and Multiplication

Cases 20, 21, 22, and 23 in Section 6.1 all involve differentiation of velocities to determine the strain rate tensor \mathbf{S} , rotation rate tensor \mathbf{R} , and their gradients $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$. Errors introduced by discrete derivative approximations involved in obtaining these quantities would then be conflated with errors that result from autonomic closure based on these frame-invariant representations. This could thereby potentially obscure the insights being sought in this Chapter regarding the fundamental accuracy that is attainable when these frame-invariant generalized representations are used in autonomic closure. For this reason, in this study all derivatives were computed in the spectral domain, where multiplication by ik_j allows exact evaluation of the derivative $\partial/\partial x_j$. This spectral differentiation ensured that \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$, and $\nabla\mathbf{R}$ were all obtained without *any* differentiation error.

Additionally, Cases 10, 11, 20, 21, 22, and 23 in Section 6.1 all involve products of velocities or velocity gradients. For Cases 10, 11, 20 and 23 these products extend only up to 2nd-order, but for Cases 21 and 22 they extend up to 5th-order. Aliasing errors introduced by each order of multiplication, especially for the higher-order products, would thus be conflated with errors that result from autonomic closure based on these frame-invariant representations. As with differentiation errors, such aliasing errors could thereby potentially obscure the insights being sought in this Chapter regarding the fundamental accuracy that is attainable when these frame-invariant generalized representations are used in autonomic closure. For this reason, in this study all such products were evaluated by first zero-padding each field in the spectral domain to the extent needed so that, when these fields were then inverse Fourier-transformed into the physical domain and multiplied to evaluate these products, the resulting products were completely free of any aliasing errors. Thus for any product of order n , for each field involved in that product this required zero-padding the field in the spectral domain by a factor of 2^{n-1} . For example, for 5th-order products the spectral domain was extended along each k_i direction by a factor of 16 and zero-padded before transforming the fields into the physical domain for multiplication. The resulting complete alias-free multiplication ensured that all terms in each of these frame-invariant representations could be evaluated without *any* aliasing errors.

The combined use of spectral differentiation and fully alias-free multiplication in principle allows the present study to unequivocally determine the fundamental accuracy that is attainable when these frame-invariant generalized representations are used in autonomic closure.

6.3 Specific Implementations of Invariant Representations

Table 2 lists the relevant parameters discussed in Section 2.4 for each implementation of autonomic closure based on the frame-invariant representations considered in this chapter. These include the number of degrees-of-freedom N , the bounding box size n^3 , the number M of training point locations, the training point spacing relative to the test-filter length scale, the effective number $M' = 6M$ of training data sets, the effective number M'/N of training data sets per degree-of-freedom, and the single-core computational time needed for autonomic closure to

Table 2. Frame-invariant implementations of autonomic closure considered in Chapter 5, showing code, number N of degrees-of-freedom in F_{ij} , relative bounding box size n^3 , number M of training points in bounding box, relative training point spacing, number M' of training data in bounding box, number of training data per degree-of-freedom M'/N , and computational time

Case	Code	N	$n^3 = \frac{V_B}{\Delta^3}$	M	$\frac{(V_B/M)^{1/3}}{\Delta}$	M'	M'/N	t(s)
10	Smith-UN	379	7^3	343	1.0	2058	5.4	2591
11	Smith-UC	28	3^3	27	1.0	162	5.8	4.0
20	Smith-19	20	3^3	27	1.0	162	8.3	5.6
21	Pope-11	11	3^3	27	1.0	162	14.7	3.8
22	ASU-8	8	3^3	27	1.0	162	20.3	2.8
23	Smith-5	5	3^3	27	1.0	162	32.4	2.2

determine the subgrid stress field. The cases are listed in order of decreasing number N of degrees-of-freedom in their underlying generalized representation F_{ij} .

Note that all these implementations have $M'/N \geq \mathcal{O}(4)$, which was found in Chapter 4 to provide sufficient training that further increases in the effective number M' of training data sets provided no added benefit. Note also that, since these frame-invariant representations involve just $1 \times N$ coefficients rather than $6 \times N$ as was the case for the *ad hoc* series representations in Chapter 4, this together with the smaller number N of degrees-of-freedom in these representations allows the use of far smaller bounding boxes than was the case in Chapter 4. That in turn allows a far more *local* determination of the coefficients, which thereby are more able to reflect the local turbulence state at the center of the bounding box where the stresses are being evaluated.




Figure 25. Scaled timings of frame-invariant implementations

6.4 Resulting Computational Times

The single-core computational time for each of these frame-invariant implementations of autonomic closure is given in Section 6.3. These computational times are compared in Figure 25 with the $N^3[1 + M/N]$ scaling identified in Section 3.6. This scaling again agrees well with the observed computational times, though as might be expected the proportionality constant is different for these frame-invariant representations than for the *ad hoc* series-based representations in Chapter 4. As expected Cases 20-23, which involve calculating \mathbf{S} and \mathbf{R} and are indicated by the red symbols in Figure 25, show a slightly different scaling than do the velocity-based representations indicated by the black symbols, which require no such overhead calculations of these tensor quantities. This is especially true for Case 20, which additionally requires calculating products of $\nabla \mathbf{S}$ and $\nabla \mathbf{R}$.

6.5 Resulting Performance of Frame-Invariant Representations

Figures 26-28 compare typical results from *a priori* tests, analogous to those described in Section 3.1, for each the frame-invariant representations in Table 2. These figures also show results obtained with the recommended series-based representation from Chapter 4 (Case 3a in Table 1, termed “CL24”), as well as the corresponding true fields to allow comparisons. Figures

26 shows results for a typical normal stress component field $\tau_{11}(\mathbf{x}, t)$, Figure 27 shows results for a typical corresponding shear stress component field $\tau_{12}(\mathbf{x}, t)$, and Figure 28 shows results for the corresponding subgrid production field $P(\mathbf{x}, t)$. In each figure, the top left panel is the true field, and the remaining panels are grouped into the velocity-based representations (CL24 and Cases 10 and 11) in the upper half of the figure, and the tensor-based representations (Cases 20-23) in the lower half of the figure.

In comparing among these results, it is immediately apparent that the velocity-based representations in the upper half of each figure agree far better with the true field than do any of the tensor-based representations in the lower half of each figure. Among the three velocity-based representations, the 6×244 -coefficient *ad hoc* series representation “CL24” in the upper right panel most accurately reproduces the true fields, but it is computationally burdensome. Of particular interest are the results from the frame-invariant velocity-based representation “Smith-UC” (Case 11), which retains only the colocated velocity products in Equation (5.1) and thus involves just 1×28 coefficients. It can be seen to reproduce the true fields nearly as accurately as does CL24, but at a dramatically lower computational cost. Furthermore, in agreement with the results from Chapter 4, “Smith-UN” (Case 10) which includes the non-colocated velocity products, performs no better than “Smith-UC”.

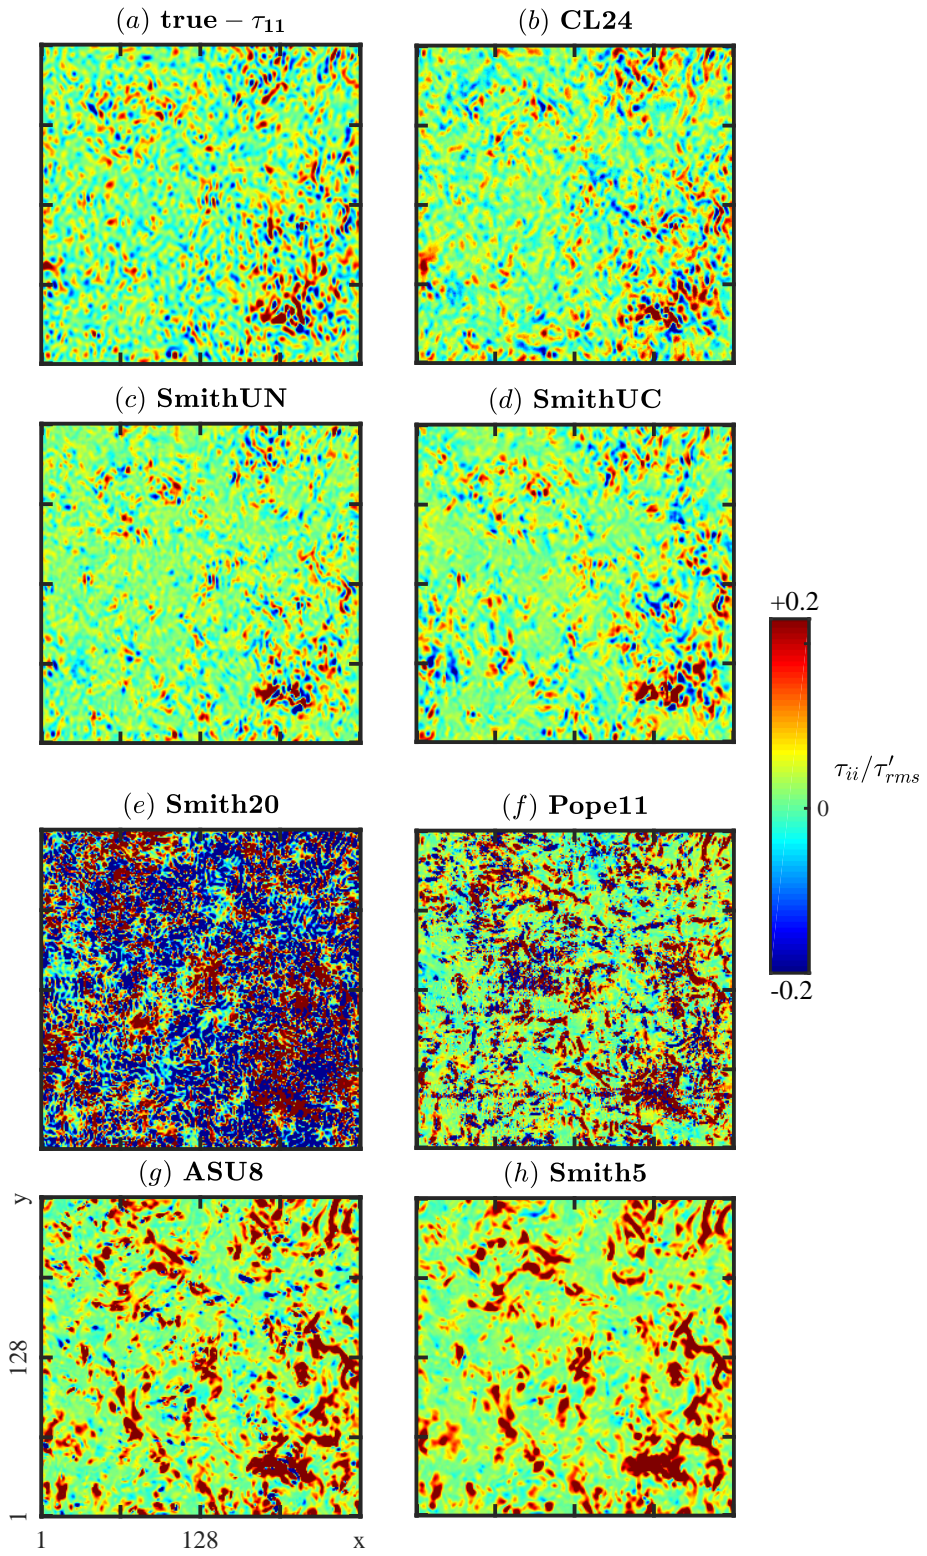


Figure 26. Comparison of (a) true stress fields τ_{11} and those computed by (b) CL24, (c) Smith-UN, (d) Smith-UC, (e) Smith-20, (f) Pope-11, (g) ASU-8, and (h) Smith-5.

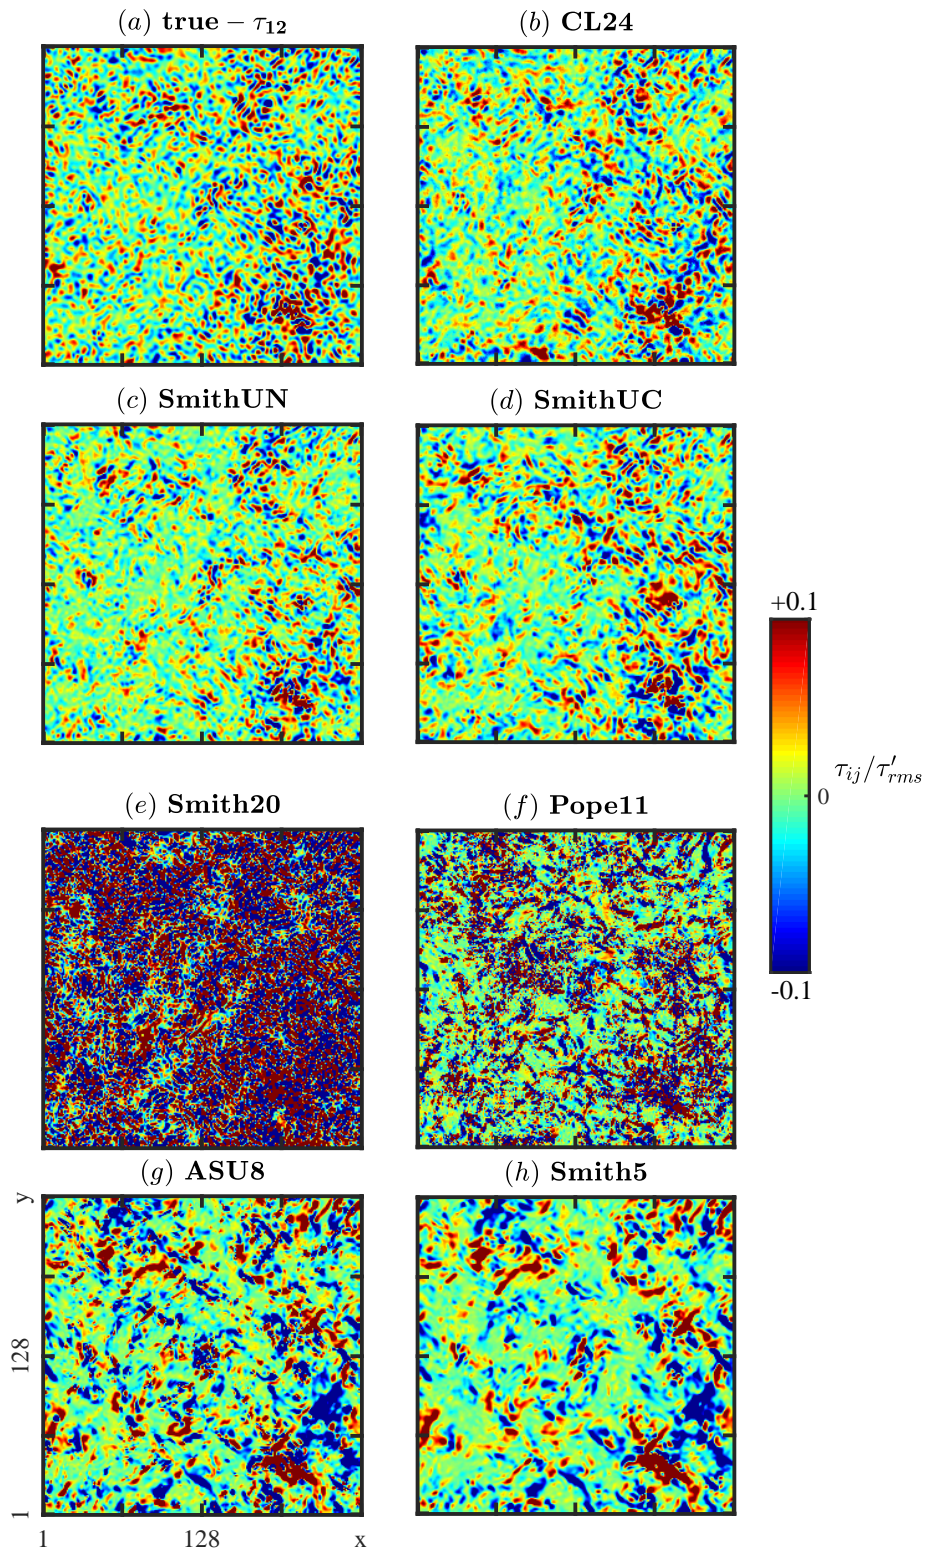


Figure 27. Comparison of (a) true stress fields τ_{12} and those computed by (b) CL24, (c) Smith-UN, (d) Smith-UC, (e) Smith-20, (f) Pope-11, (g) ASU-8, and (h) Smith-5.

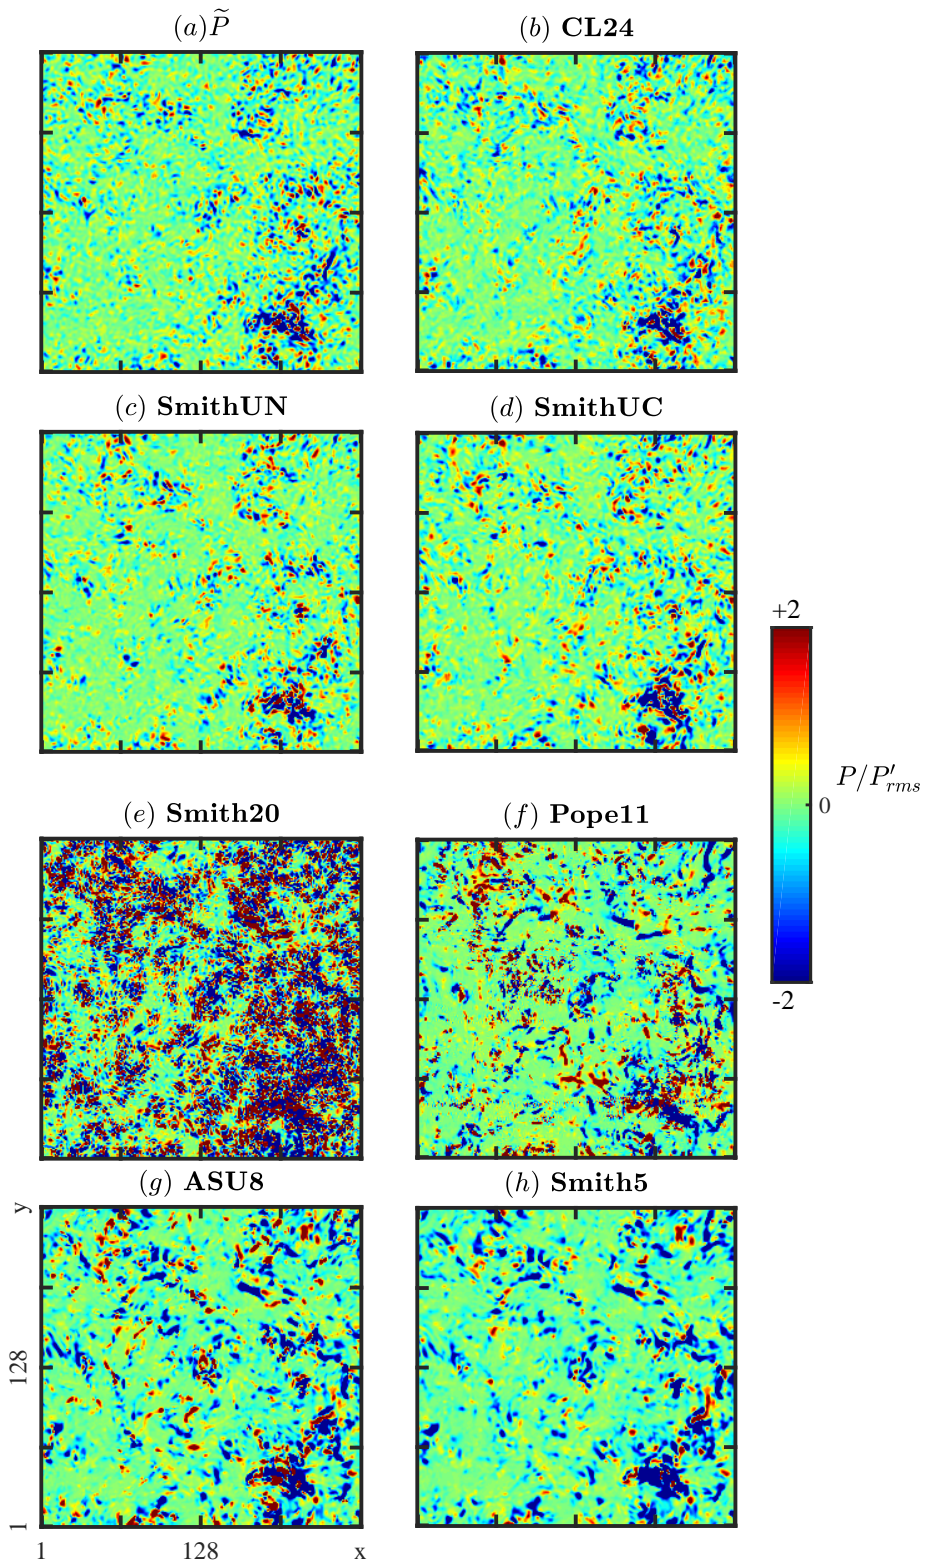


Figure 28. Comparison of (a) true production fields and those computed by (b) CL24, (c) Smith-UN, (d) Smith-UC, (e) Smith-20, (f) Pope-11, (g) ASU-8, and (h) Smith-5.




Figure 29. Comparison of PDFs for (a) stress field τ_{12} , and (b) production field P

The tensor-based frame-invariant representations in the lower half of each of these figures (Cases 20-23, which involve \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$) can be seen to reproduce the true fields far less accurately than do any of the velocity-based representations in the upper half. This was surprising, since many of these representations retain higher-order products of the stencil velocities, via the various tensor contractions involved in their underlying representations, than do the second-order velocity products in “CL24” (Case 3a), “Smith-UN” (Case 10), and “Smith-UC” (Case 11).

The reasons for this cannot be due to differentiation errors introduced by the required velocity gradients, since as noted in Section 6.2 spectral differentiation was used in these tests precisely to avoid any such differentiation errors. The poor performance of Cases 20-23 also cannot be due to aliasing errors introduced by the products (which range from 2nd-order to 5th-order) involved in these representations, since as noted in Section 6.2 extensive zero-padding was used to completely avoid *any* aliasing in *any* of these products. For these reasons it was surprising that the tensor-based frame-invariant representations in the lower half of Figures 26-28 performed as poorly as can be seen in those results.

6.6 Eigenvalue Distributions in the Generalized Representations

This section examines the linear algebra involved in the matrix inversion that is used to solve the optimization problem for the coefficients in each of the generalized representations

F_{ij} , namely each of the cases listed in Table 1 and Table 2. The damped least-squares optimization used throughout this work obtains the coefficients from the test stresses T_{ij} via an inverse solution of the general form

$$\mathbf{h}_{ij} = \left(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I} \right)^{-1} \widehat{\mathbf{V}}^T \mathbf{T}_{ij}. \quad (6.7)$$

It is only the $\widehat{\mathbf{V}}$ matrix that differs from one generalized representation to another. The different representations thereby produce different structures in this matrix, which in turn produce widely varying differences in the structure of the corresponding $(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I})$ matrix, the inverse of which must be computed.

This can be seen in Figures 30-36, which show these matrix structures and the corresponding eigenvalue distributions for the three velocity-based representations in CL24, Smith-UC, and Smith-UN, as well as the four tensor-based representations in Smith-5, ASU-8, Pope-11, and Smith-19. In each figure, the upper left panel shows the logarithmic magnitude of the entries in the $M \times N$ dimensional \mathbf{V} matrix, and the upper right panel shows the logarithmic magnitudes in the corresponding $(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I})$ matrix. Comparing the range of magnitudes in the $(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I})$ matrix in the upper right panel of each figure, it is apparent that the three velocity-based representations have a much narrow range of magnitudes than do the tensor-based representations. For the tensor-based representations in Figures 33-36, these magnitudes typically span 18-20 orders of magnitude (though for the ASU-8 representation the span is only about 5 orders of magnitude). In contrast, for the velocity-based representations in Figures 30-32 the magnitudes span a far narrower range, making accurate inversion more readily possible. Notably the span is by far smallest for the Smith-UC representation (Case 11), which was seen in Figures 26-28 to perform very well.

These observations are most directly reflected in the lower part of each of Figures 30-36, which show the eigenvalue distributions for the $(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I})$ matrix that must be inverted in each of these representations. In particular, for the tensor-based representations the eigenvalues typically span from $10^2 - 10^5$, and in the case of Smith-19 in Figure 36 they are as large as 10^8 . In all of these cases, the eigenvalues are far removed from $\mathcal{O}(1)$, indicating that accurate

inversion of the $(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I})$ would at a minimum require preconditioning, and may be difficult to obtain even with preconditioning.

In contrast, the eigenvalue distributions in Figures 30-32 for the velocity-based representations show that these are more readily able to allow accurate inversion, since the spectral radii of most of the eigenvalues satisfy $|\lambda_i| \leq 1$. In particular, for the 28-term Smith-UC (Case 11) representation the eigenvalues are generally distributed within $10^{-1} \leq \lambda_i \leq 10^1$, which should allow accurate inversion even without extensive preconditioning.

In the present study, standard LAPACK routines were used to obtain the inverse $(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I})^{-1}$, which make use of basic preconditioning methods but do not include more recently developed advanced preconditioning methods of the type that are beginning to be introduced in linear algebra packages from Intel and others. The results in this section and those in Section 6.5 suggest that the limited accuracy obtained with the tensor-based frame-invariant representations developed and tested in this study is due to these limitations in the accuracy with which the inverse can be computed. Indeed, the trends seen in the relative accuracy of the results in Figures 26-28 is consistent with the observation in this section with regard to the suitability of their respective $(\widehat{\mathbf{V}}^T \widehat{\mathbf{V}} + \lambda \mathbf{I})$ matrices for accurate inversion.




Figure 30 (a). Typical linear system in CL24 represented by (a) V matrix (left) (b) $V^T V + \lambda I$ matrix (right)



Figure 30 (b). Eigenvalue (λ_i) distribution for CL24.

Figure 30. Inverse system for CL24.




Figure 31 (a). Typical linear system in SmithUC represented by (a) V matrix (left) (b) $V^T V + \lambda I$ matrix (right)



Figure 31 (b). Eigenvalue (λ_i) distribution for SmithUC.

Figure 31. Inverse system for SmithUC.

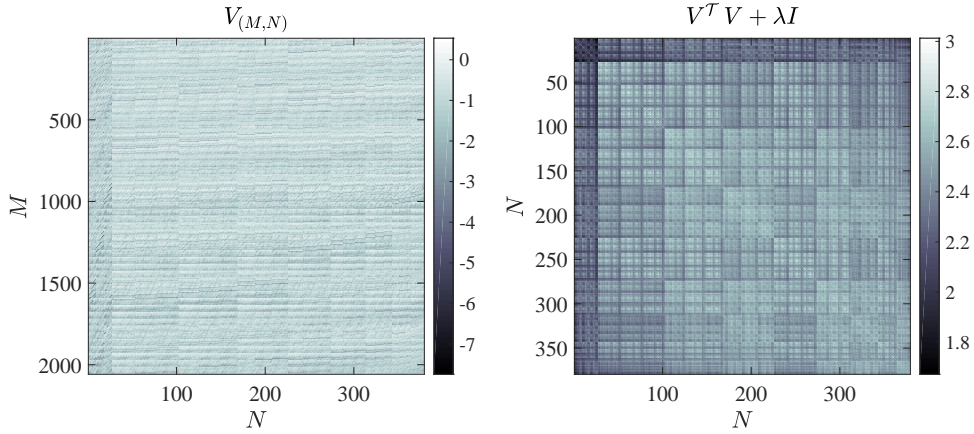


Figure 32 (a). Typical linear system in SmithUN represented by (a) V matrix (left) (b) $V^T V + \lambda I$ matrix (right)




Figure 32 (b). Eigenvalue (λ_i) distribution for SmithUN.

Figure 32. Inverse system for SmithUN.




Figure 33 (a). Typical linear system in Smith-5 represented by (a) V matrix (left) (b) $V^T V + \lambda I$ matrix (right)




Figure 33 (b). Eigenvalue (λ_i) distribution for Smith-5.

Figure 33. Inverse system for Smith5.




Figure 34 (a). Typical linear system in ASU8 represented by (a) V matrix (left) (b) $V^T V + \lambda I$ matrix (right)




Figure 34 (b). Eigenvalue (λ_i) distribution for ASU8.

Figure 34. Inverse system for ASU8.




Figure 35 (a). Typical linear system in Pope11 represented by (a) V matrix (left) (b) $V^T V + \lambda I$ matrix (right)




Figure 35 (b). Eigenvalue (λ_i) distribution for Pope11.

Figure 35. Inverse system for Pope11.




Figure 36 (a). Typical linear system in Smith19 represented by (a) V matrix (left) (b) $V^T V + \lambda I$ matrix (right)




Figure 36 (b). Eigenvalue (λ_i) distribution for Smith19.

Figure 36. Inverse system for Smith19.

This leaves open the possibility that more advanced preconditioning methods could allow highly efficient representations such as ASU-8 (Case 22) and even Smith-5 (Case 23), which respectively involve just 5 and 8 coefficients, to be accurately implemented in autonomic closure. Investigation of such advanced preconditioning methods was beyond the scope of this study, but is an obvious area of research that could provide even greater accuracy and efficiency in autonomic closure.

6.7 Recommended Frame-Invariant Representation

Based on the results in Chapter 4 and those in this Chapter 6, the 1×28 -coefficient velocity-based frame-invariant representation in Equation (6.2) is recommended as the most accurate and efficient representation for the subgrid stress in autonomic closure. Due to the small number of coefficients involved, its efficiency is such that it can be implemented in a fully dynamic “on-the-fly” manner in large eddy simulations that are based on autonomic closure. Based on these results, it can be expected that for other subgrid terms in large eddy simulations, similar frame-invariant generalized representations expressed directly in the primitive variables of the simulation may provide comparable accuracy and efficiency.

CONCLUSIONS AND IMPLICATIONS FOR LES

7.1 Summary

Autonomic closure, originally proposed in Refs. [58, 59], is a new general methodology for subgrid closures that circumvents the need to specify a particular fixed closure model, and instead allows a fully-adaptive self-optimizing closure methodology. The closure is autonomic in the sense that the simulation itself determines the optimal relation at each point and time between any subgrid term and the primitive variables in the simulation, through the solution of a local system identification problem. It is based on highly generalized representations of subgrid terms, and solution of a local system identification problem to find the representation that best connects the local subgrid term to the local primitive variables. These generalized representations have a substantial number of degrees of freedom that are determined dynamically at each point \mathbf{x} and time t in the simulation. This eliminates the need for a traditional predefined closure model, and allows autonomic closure to freely adapt to widely varying degrees of nonlinearity, nonlocality, nonequilibrium, and other characteristics in the turbulence state at each point and time in the simulation. Importantly, autonomic closure involves no explicit external length scale, such as the LES or test filter length scales that appear in most traditional prescribed closure models such as the widely-used basic and dynamic versions of the Smagorinsky model and nearly all other subgrid models. Indeed autonomic closure is not a closure model; instead it is a closure methodology that enables essentially model-free closure of any subgrid term.

Autonomic closure can be regarded as a high-dimensional generalization of the dynamic approach used with various traditional prescribed closure models. It can also be regarded as a type of “data-driven” turbulence closure, in which machine-learning methods are used with available data to discover a closure model rather than prescribe one. However unlike other data-

driven models, which use external training data obtained from prior simulations or experiments, autonomic closure is based on internal training data obtained at each point and time at a test-filter scale in the simulation. It then generalizes the resulting local generalized representation, in a machine learning sense, from the test-filter scale to the LES scale.

In this study, *a priori* tests were performed to develop accurate and efficient implementations of autonomic closure based on choices for the generalized representations F_{ij} and parameters associated with the local system identification of the turbulence state. These included the relative number of training points and the bounding box size, which can impact the generalizability of coefficients in the representation from the test scale to the LES scale, and also affect the computational cost of this closure methodology. The main focus has been on studying the impacts of these factors on the resulting accuracy and efficiency of autonomic closure for the subgrid stress $\tau_{ij}(\mathbf{x}, t)$. Particular attention has been paid in this study to the associated subgrid production $P(\mathbf{x}, t) = -\tau_{ij}\widetilde{S}_{ij}$, including the extent to which autonomic closure can correctly produce the same *structural* features in which large magnitudes of forward and backward scatter concentrated in the true $P(\mathbf{x}, t)$ fields.

7.2 Major Conclusions from This Study

1. A reduction of more than five orders of magnitude in the computational cost of autonomic closure was obtained in this study compared to Refs. [58, 59], with essentially no loss of accuracy in the resulting subgrid stress and production fields. This was primarily achieved by using more efficient generalized representations, including forms that leverage the frame-invariance properties of the subgrid stress to greatly reduce the number of degrees of freedom in this representation.
2. The general formulation of Smith (1971) for a *complete and minimal* tensor basis in any polynomial representation of any symmetric rank-two tensor in terms of any number of vectors and tensors was used in this study to obtain highly compact generalized representations for use in autonomic closure. This included a general representation for

the subgrid stress in terms of \mathbf{S} , \mathbf{R} , $\nabla\mathbf{S}$ and $\nabla\mathbf{R}$. For a representation involving only \mathbf{S} , \mathbf{R} , the resulting tensor basis is smaller than the previously assumed minimal basis proposed by Lumley (1970) and Pope (1974).

3. The subgrid production support-density field $G(\mathbf{x}, t)$ was developed as a means for quantifying where highly intermittent large positive and negative values of the subgrid production are concentrated across all scales of the simulation, and metrics \mathcal{M}_1 and \mathcal{M}_2 were developed to allow quantitative assessments of the structural similarity in two such highly intermittent fields. These metrics can be used in *a priori* tests to effectively compare the detailed spatial structure of the subgrid production from any closure model with the true production field.
4. Even when using the highly efficient generalized representations F_{ij} developed in this study, the subgrid production fields $P^F(\mathbf{x}, t)$ obtained from autonomic closure are seen to be accurate in nearly all details of their spatial structure when compared with the corresponding true fields $P(\mathbf{x}, t)$.
5. The accuracy of the subgrid stress fields $\tau_{ij}^F(\mathbf{x}, t)$ and the associated subgrid production fields $P^F(\mathbf{x}, t)$ obtained from these accurate and efficient implementations of autonomic closure far exceeds the accuracy with which these fields are represented by traditional prescribed closure models, including the dynamic Smagorinsky model and the Bardina model. This applies also to subgrid stress and production statistics obtained from these fields.
6. There are no assumptions in the autonomic closure methodology that depend in any way on the local state of the turbulence or the particular type of turbulent flow in which this closure methodology is applied. The only requirement is that the relation between the test stresses T_{ij} and the test-filtered variables in (2.1) must be the same as that between the subgrid stresses τ_{ij} and the LES-filtered variables in (2.2); this should be valid when both the test-filter scale and the LES-filter scale are in the scale-similar inertial range.

7. Consequently, autonomic closure was found in this study to perform equally well in tests based on highly sheared homogeneous turbulence as it did in tests with isotropic turbulence.
8. Based on the results obtained in Chapters 4 and 6, the recommended form for the generalized representation of the subgrid stress in autonomic closure is the 28-coefficient frame-invariant velocity-based formulation denoted as Case 11 (Smith-UC) in Table 2. This form for F_{ij} , and the small bounding box sizes that it enables, provides subgrid stress and production fields that are essentially as accurate as far more complex and computationally costly representations.
9. Even more compact generalized representations, including Cases 22 and 23 in Table 2, which respectively involve just 8 and 5 coefficients, may allow even greater accuracy and efficiency, though their spectral radii require levels of preconditioning that were beyond the scope of this study. While this is an obvious area for future research, the present study has identified representations that are sufficiently accurate and efficient for practical multi-physics large eddy simulations, where the accuracy provided by autonomic closure across essentially all resolved scales can be particularly important.

REFERENCES

- [1] J. C. Tyacke, M. Mahak, and P. G. Tucker. “Large-scale multifidelity, multiphysics, hybrid Reynolds-averaged Navier–Stokes/large-eddy simulation of an installed aero-engine”. *J. Prop. Power* 32.4 (2016), pp. 997–1008.
- [2] P. Schmitt, T. Poinso, B. Schuermans, and K. Geigle. “Large-eddy simulation and experimental study of heat transfer, nitric oxide emissions and combustion instability in a swirled turbulent high-pressure burner”. *J. Fluid Mech.* 570 (2007), pp. 17–46.
- [3] A. Neophytou, B. Cuenot, and P. Duchaine. “Large-eddy simulation of ignition and flame propagation in a trisector combustor”. *J. Prop. Power* 32.2 (2015), pp. 345–359.
- [4] R. Baurle. “Hybrid Reynolds-averaged/large-eddy simulation of a scramjet cavity flameholder”. *AIAA J.* 55.2 (2016), pp. 544–560.
- [5] A. Poubeau, R. Paoli, A. Dauplain, F. Duchaine, et al. “Large-eddy simulations of a single-species solid rocket booster jet”. *AIAA J.* 53.6 (2014), pp. 1477–1491.
- [6] F. Xiao, M. Dianat, and J. J. McGuirk. “Large eddy simulation of liquid-jet primary breakup in air crossflow”. *AIAA J.* 51.12 (2013), pp. 2878–2893.
- [7] M. Bini and W. Jones. “Large-eddy simulation of particle-laden turbulent flows”. *J. Fluid Mech.* 614 (2008), pp. 207–252.
- [8] G. Aubard, X. Gloerfelt, and J.-C. Robinet. “Large-eddy simulation of broadband unsteadiness in a shock/boundary-layer interaction”. *AIAA J.* 51.10 (2013), pp. 2395–2409.
- [9] M. S. Loginov, N. A. Adams, and A. A. Zheltovodov. “Large-eddy simulation of shock-wave/turbulent-boundary-layer interaction”. *J. Fluid Mech.* 565 (2006), pp. 135–169.
- [10] U. Piomelli. “Large-eddy simulation: achievements and challenges”. *Prog. Aerosp. Sci.* 35.4 (1999), pp. 335–362.
- [11] K. Mahesh, G. Constantinescu, S. Apte, G. Iaccarino, et al. “Large-eddy simulation of reacting turbulent flows in complex geometries”. *J. Appl. Mech.* 73.3 (2006), pp. 374–381.
- [12] S. Bose and P. Moin. “A dynamic slip boundary condition for wall-modeled large-eddy simulation”. *Phys. Fluids* 26.1 (2014), p. 015104.
- [13] U. Piomelli. “Wall-layer models for large-eddy simulations”. *Prog. Aerosp. Sci.* 44.6 (2008), pp. 437–446.

- [14] U. Piomelli. “Wall-modeled large-eddy simulations: Present status and prospects”. In: *Direct and Large-Eddy Simulation VII*. Springer, 2010, pp. 1–10.
- [15] S. Kawai and J. Larsson. “Wall-modeling in large eddy simulation: Length scales, grid resolution, and accuracy”. *Phys. Fluids* 24.1 (2012), p. 015105.
- [16] M. Dianat, Z. Yang, D. Jiang, and J. McGuirk. “Large eddy simulation of scalar mixing in a coaxial confined jet”. *Flow Turbul. Combust.* 77.1-4 (2006), p. 205.
- [17] G. C. Burton. “The nonlinear large-eddy simulation method applied to $Sc \approx 1$ and $Sc \gg 1$ passive-scalar mixing”. *Phys. Fluids* 20.3 (2008), p. 035103.
- [18] G. C. Burton. “Study of ultrahigh Atwood-number Rayleigh–Taylor mixing dynamics using the nonlinear large-eddy simulation method”. *Phys. Fluids* 23.4 (2011), p. 045106.
- [19] J. M. Mejía, A. Sadiki, A. Molina, F. Chejne, et al. “Large eddy simulation of the mixing of a passive scalar in a high-Schmidt turbulent jet”. *J. Fluids Eng.* 137.3 (2015), p. 031301.
- [20] S. Elghobashi. “On predicting particle-laden turbulent flows”. *Appl. Sci. Res.* 52.4 (1994), pp. 309–329.
- [21] F. Xiao, Z. Wang, M. Sun, J. Liang, et al. “Large eddy simulation of liquid jet primary breakup in supersonic air crossflow”. *Int. J. Multiphase Flow* 87 (2016), pp. 229–240.
- [22] A. Irannejad and F. Jaber. “Large eddy simulation of turbulent spray breakup and evaporation”. *Int. J. Multiphase Flow* 61 (2014), pp. 108–128.
- [23] S. De and S. H. Kim. “Large eddy simulation of dilute reacting sprays: Droplet evaporation and scalar mixing”. *Comb. Flame* 160.10 (2013), pp. 2048–2066.
- [24] H. Pitsch. “Large-eddy simulation of turbulent combustion”. *Annu. Rev. Fluid Mech.* 38 (2006), pp. 453–482.
- [25] B. Franchetti, F. C. Marincola, S. Navarro-Martinez, and A. Kempf. “Large eddy simulation of a pulverised coal jet flame”. *Proc. Comb. Inst.* 34.2 (2013), pp. 2419–2426.
- [26] D. K. Tafti, L. He, and K. Nagendra. “Large eddy simulation for predicting turbulent heat transfer in gas turbines”. *Phil. Trans. R. Soc. A* 372.2022 (2014), p. 20130322.
- [27] C. Meneveau and J. Katz. “Scale-invariance and turbulence models for large-eddy simulation”. *Annu. Rev. Fluid Mech.* 32.1 (2000), pp. 1–32.
- [28] S. Liu, C. Meneveau, and J. Katz. “On the properties of similarity subgrid-scale models as deduced from measurements in a turbulent jet”. *J. Fluid Mech.* 275 (1994), pp. 83–119.

- [29] P. Sagaut. *Large eddy simulation for incompressible flows: An Introduction*. Springer Berlin, 2006.
- [30] S. Menon, P.-K. Yeung, and W.-W. Kim. “Effect of subgrid models on the computed interscale energy transfer in isotropic turbulence”. *Comput. Fluids* 25.2 (1996), pp. 165–180.
- [31] Y. Shi, Z. Xiao, and S. Chen. “Constrained subgrid-scale stress model for large eddy simulation”. *Phys. Fluids* 20.1 (2008), p. 011701.
- [32] A. E. Tejada-Martinez and K. E. Jansen. “A dynamic Smagorinsky model with dynamic determination of the filter width ratio”. *Phys. Fluids* 16.7 (2004), pp. 2514–2528.
- [33] C. Yu, Z. Xiao, and X. Li. “Scale-adaptive subgrid-scale modelling for large-eddy simulation of turbulent flows”. *Phys. Fluids* 29.3 (2017), p. 035101.
- [34] C. B. da Silva and J. C. Pereira. “The effect of subgrid-scale models on the vortices computed from large-eddy simulations”. *Phys. Fluids* 16.12 (2004), pp. 4506–4534.
- [35] S. Ghosal. “An analysis of numerical errors in large-eddy simulations of turbulence”. *J. Comput. Phys.* 125.1 (1996), pp. 187–206.
- [36] J. Meyers, B. J. Geurts, and M. Baelmans. “Database analysis of errors in large-eddy simulation”. *Phys. Fluids* 15.9 (2003), pp. 2740–2755.
- [37] A. Kravchenko and P. Moin. “On the effect of numerical errors in large eddy simulations of turbulent flows”. *J. Comput. Phys.* 131.2 (1997), pp. 310–322.
- [38] F. K. Chow and P. Moin. “A further study of numerical errors in large-eddy simulations”. *J. Comput. Phys.* 184.2 (2003), pp. 366–380.
- [39] M. Germano. “Turbulence: the filtering approach”. *J. Fluid Mech.* 238 (1992), pp. 325–336.
- [40] S. Pope. *Turbulent Flows*. Cambridge University Press, Cambridge, UK, 2000, pp. 1361–1362.
- [41] T. Lund and H. Kaltenbach. *Experiments with explicit filtering for LES using a finite-difference method*. Annual Research Briefs, 1995, pp. 91–105. Center for Turbulence Research, Stanford University.
- [42] J. Smagorinsky. “General circulation experiments with the primitive equations”. *Mon. Weather Rev.* 91 (1963), p. 99.
- [43] J. Bardina, J. Ferziger, and W. Reynolds. “Improved subgrid-scale models for large-eddy simulation”. *AIAA* 80-1357 (1980).

- [44] J. Bardina, J. H. Ferziger, and W. C. Reynolds. *Improved turbulence models based on large eddy simulation of homogeneous, incompressible turbulent flows*. Tech. rep. TF-19. Stanford University, 1983.
- [45] K. Horiuti. “The role of the Bardina model in large eddy simulation of turbulent channel flow”. *Phys. Fluids A* 1.2 (1989), pp. 426–428.
- [46] M. Germano, U. Piomelli, P. Moin, and W. H. Cabot. “A dynamic subgrid-scale eddy viscosity model”. *Phys. Fluids A* 3.7 (1991), pp. 1760–1765.
- [47] D. K. Lilly. “A proposed modification of the Germano subgrid-scale closure method”. *Phys. Fluids A* 4.3 (1992), pp. 633–635.
- [48] P. Moin, K. Squires, W. Cabot, and S. Lee. “A dynamic subgrid-scale model for compressible turbulence and scalar transport”. *Phys. Fluids A* 3.11 (1991), pp. 2746–2757.
- [49] U. Piomelli. “High Reynolds number calculations using the dynamic subgrid-scale stress model”. *Phys. Fluids A* 5.6 (1993), pp. 1484–1490.
- [50] Y. Zang, R. L. Street, and J. R. Koseff. “A dynamic mixed subgrid-scale model and its application to turbulent recirculating flows”. *Phys. Fluids A* 5.12 (1993), pp. 3186–3196.
- [51] R. Akhavan, A. Ansari, S. Kang, and N. Mangiavacchi. “Subgrid-scale interactions in a numerically simulated planar turbulent jet and implications for modelling”. *J. Fluid Mech.* 408 (2000), pp. 83–120.
- [52] S. Liu, C. Meneveau, and J. Katz. “Experimental study of similarity subgrid-scale models of turbulence in the far-field of a jet”. *Appl. Sci. Res.* 54.3 (1995), pp. 177–190.
- [53] S. Liu, J. Katz, and C. Meneveau. “Evolution and modelling of subgrid scales during rapid straining of turbulence”. *J. Fluid Mech.* 387 (1999), pp. 281–320.
- [54] F. Sarghini, U. Piomelli, and E. Balaras. “Scale-similar models for large-eddy simulations”. *Phys. Fluids* 11.6 (1999), pp. 1596–1607.
- [55] T. S. Lund, S. Ghosal, and P. Moin. “Numerical experiments with highly-variable eddy viscosity models”. *Engineering Applications to Large Eddy Simulation*. Ed. by R. S. Piomelli U. Vol. 162. New York: ASME, 1993, pp. 7–11.
- [56] S. Ghosal, T. S. Lund, P. Moin, and K. Akselvoll. “A dynamic localization model for large-eddy simulation of turbulent flows”. *J. Fluid Mech.* 286 (1995), pp. 229–255.
- [57] V. Gravemeier. “A consistent dynamic localization model for large eddy simulation of turbulent flows based on a variational formulation”. *J. Comput. Phys.* 218.2 (2006), pp. 677–701.

- [58] R. N. King, P. E. Hamlington, and W. Dahm. “Autonomic subgrid-scale closure for large eddy simulations”. *AIAA Paper* 2015-1285 (2015).
- [59] R. N. King, P. E. Hamlington, and W. J. Dahm. “Autonomic closure for turbulence simulations”. *Phys. Rev. E* 93.3 (2016), p. 031301.
- [60] U. Piomelli, W. H. Cabot, P. Moin, and S. Lee. “Subgrid-scale backscatter in turbulent and transitional flows”. *Phys. Fluids A* 3.7 (1991), pp. 1766–1771.
- [61] P. E. Hamlington and W. J. Dahm. “Reynolds stress closure for nonequilibrium effects in turbulent flows”. *Phys. Fluids* 20.11 (2008), p. 115101.
- [62] P. E. Hamlington and W. J. Dahm. “Nonlocal form of the rapid pressure-strain correlation in turbulent flows”. *Phys. Rev. E* 80.4 (2009), p. 046311.
- [63] Z. J. Zhang and K. Duraisamy. “Machine learning methods for data-driven turbulence modeling”. *AIAA* 2015-2460 (2015).
- [64] K. Duraisamy. “Informing turbulence closures with computational and experimental data”. *AIAA* 2016-1556 (2016).
- [65] A. P. Singh, K. Duraisamy, and Z. J. Zhang. “Augmentation of turbulence models using field inversion and machine learning”. *AIAA* 2017-0993 (2017).
- [66] J. Ling and J. Templeton. “Evaluation of machine learning algorithms for prediction of regions of high Reynolds averaged Navier Stokes uncertainty”. *Phys. Fluids* 27.8 (2015), p. 085103.
- [67] J. Ling, A. Kurzawski, and J. Templeton. “Reynolds averaged turbulence modelling using deep neural networks with embedded invariance”. *J. Fluid Mech.* 807 (2016), pp. 155–166.
- [68] J. Ling, R. Jones, and J. Templeton. “Machine learning strategies for systems with invariance properties”. *J. Comput. Phys.* 318 (2016), pp. 22–35.
- [69] J.-X. Wang and H. Xiao. “Data-driven CFD modeling of turbulent flows through complex structures”. *Int. J. Heat and Fluid.* 62 (2016), pp. 138–149.
- [70] J.-X. Wang, J.-L. Wu, and H. Xiao. “Physics-informed machine learning approach for reconstructing Reynolds stress modeling discrepancies based on DNS data”. *Phys. Rev. Fluids* 2.3 (2017), p. 034603.
- [71] J.-L. Wu, J.-X. Wang, H. Xiao, and J. Ling. “A priori assessment of prediction confidence for data-driven turbulence modeling”. *Flow Turbul. Combust.* 99.1 (2017), pp. 25–46.

- [72] P. SB. “A more general effective-viscosity hypothesis”. *J. Fluid Mech.* 72.2 (1975), pp. 331–340.
- [73] M. Oberlack. *Invariant modeling in large-eddy simulation of turbulence*. Annual Research Briefs, 1997, pp.91-105. Center for Turbulence Research, Stanford University.
- [74] D. Razafindralandy, A. Hamdouni, and M. Oberlack. “Analysis and development of subgrid turbulence models preserving the symmetry properties of the Navier–Stokes equations”. *Eur. J. Mech. B/Fluids* 26.4 (2007), pp. 531–550.
- [75] M. Schetzen. *The Volterra and Wiener Theories of Nonlinear Systems*. Melbourne, FL, USA: Kriger Publishing Co. Inc., 2006.
- [76] *Johns Hopkins University Turbulence Database*. URL: <http://turbulence.pha.jhu.edu/>.
- [77] Y. Li, E. Perlman, M. Wan, Y. Yang, et al. “A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence”. *J. Turb.* 9 (2008), N31.
- [78] A. Kshitiij, E. Stallcup, W. J. Dahm, and P. Hamlington. “Complete frame-invariant tensor formulation of autonomic closure.” White paper, in preparation.
- [79] A. Kshitiij, E. Stallcup, W. Dahm, and P. Hamlington. “Complete frame-invariant representation of the subgrid scalar flux.” White paper, in preparation.
- [80] S. Pennisi and M. Trovato. “On the irreducibility of Professor GF Smith’s representations for isotropic functions”. *Int. J. Engg. Sci.* 25.8 (1987), pp. 1059–1065.
- [81] G. Smith. “On isotropic functions of symmetric tensors, skew-symmetric tensors and vectors”. *Int. J. Engg. Sci.* 9.10 (1971), pp. 899–916.
- [82] O. Doronina, J. Christopher, C. A. Towery, P. Hamlington, et al. “Autonomic closure for turbulent flows using approximate bayesian computation”. *AIAA* 2018-0594 (2018).
- [83] T. B. Gatski and C. G. Speziale. “On explicit algebraic stress models for complex turbulent flows”. *J. Fluid Mech.* 254 (1993), pp. 59–78.
- [84] M. Itskov. *Tensor algebra and tensor analysis for engineers*. Springer, 2007.
- [85] J. Lumley. “Toward a turbulent constitutive relation”. *J. Fluid Mech.* 41.2 (1970), pp. 413–434.
- [86] T. S. Lund and E. Novikov. *Parameterization of subgrid-scale stress by the velocity gradient tensor*. Annual Research Briefs, 1992, pp. 27-43. Center for Turbulence Research, Stanford University.

- [87] A. J. M. Spencer and R. S. Rivlin. “The theory of matrix polynomials and its application to the mechanics of isotropic continua”. *Arch. Rat. Mech. Anal.* 2.1 (1958), pp. 309–336.
- [88] A. J. M. Spencer and R. S. Rivlin. “Further results in the theory of matrix polynomials”. *Arch. Rat. Mech. Anal.* 4.1 (1959), pp. 214–230.
- [89] C. Zemach. *Mathematics of Fluid Mechanics Mechanics (Appendix A in Handbook of Fluid Dynamics, R.W. Johnson, Ed.)* CRC Press, 1998.
- [90] Q.-S. Zheng. “Theory of representations for tensor functions—A unified invariant approach to constitutive equations”. *Appl. Mech. Rev.* 47.11 (1994), pp. 545–587.