

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Reference herein to any social initiative (including but not limited to Diversity, Equity, and Inclusion (DEI); Community Benefits Plans (CBP); Justice 40; etc.) is made by the Author independent of any current requirement by the United States Government and does not constitute or imply endorsement, recommendation, or support by the United States Government or any agency thereof.

Whitepaper: Optimal Control from a Fluid Dynamics Perspective

A Perloff, J Pratt, M Schneider

October 2025



Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.



Lawrence Livermore
National Laboratory

Whitepaper: Optimal Control from a Fluid Dynamics Perspective

J. Pratt , M. Schneider , A. Perloff 
Astronomy and Astrophysics Analytics Group, Physics Division

Livermore, California September 25, 2025

Abstract

An optimal control problem described by the Hamilton-Jacobi-Bellman equation can be developed into a problem that can be solved by general computational fluid dynamics packages. We describe how this formulation would allow a classical problem in optimal control, Zermelo's problem, to be treated as a multi-fluid problem. This approach has the advantage of allowing optimal navigation problems to be conducted over large areas, as well as to include moderately larger numbers of ships. We draw comparisons between this approach and the field of fluid control for fluid animations in movies.

1 | Introduction to the HJB equation for optimal control

In the field of optimal control, a control is sought so that the behavior of a dynamical system optimizes an objective function. This is most easily understood through an example. Let us consider a problem from finance [Forsyth and Labahn, 2007] where an asset S follows the stochastic PDE

$$\frac{\partial S}{\partial t} = \mu S(x, t) + \sigma S(x, t) \frac{dW}{dt} . \quad (1.1)$$

Here μ is a drift rate, and σ is the volatility, a measurement of how varied (over time) the return is on an asset. The volatility is an amplitude for the time derivative of W , the usual Wiener process. An optimal control problem for this system can be written for the value $J(S, t)$ of a claim written on asset S

$$\frac{\partial J}{\partial t} + \sup_q \left(\frac{q^2 S^2}{2} J_{SS} + S J_S - r S \right) = 0 . \quad (1.2)$$

Here r is the given borrowing rate, and a value of q is sought such that the final selling value is specified $J(S, t_f) = D$. This is an equation of the Hamilton-Jacobi-Bellman (HJB) type [e.g. as discussed by Kappen, 2005, Kühn, 2019, Bardi et al., 1997, Crespo and Sun, 2003]; the equation has the basic form

$$\frac{\partial J}{\partial t} = HJ . \quad (1.3)$$

Here the operator H is a Hamiltonian; the Hamiltonian in an HJB equation typically has a minimum or maximum involved, making these equations notoriously difficult to solve [Akian et al., 2018]. The HJB equation is typically solved backward in time, with a final time solution provided as a starting point. HJB equations appear in probability theory as the Kolmogorov backward equation [Kühn, 2019]. They thus also have a theoretical connection to the Kolmogorov forward equation, which is more commonly known as the Fokker Planck equation [Annunziato et al., 2014]. In order to solve HJB equations, viscosity solutions [Crandall, 2006] have been developed; these provide a mathematical framework for selecting a value function J from a possibly infinite set of weak solutions.

Schneider et al. [2022] examine a particular linear HJB-type equation designed to describe a linear Markov decision process (LMDP); they demonstrate that this equation can be expressed as a linear Schrödinger-type equation. That transform has also been explored with mathematical rigor by Teter et al. [2025]. In this setting, the optimal control aspect of the calculation is entirely encoded in a potential function. We refer to this potential as the Bellman potential, and also refer to the wavefunction that moves through the Bellman potential as the Bellman wavefunction. The transformation of a standard HJB equation to a Schrödinger-type equation is useful because it allows traditional scientific computing methods to be applied to optimal control. However, dealing with a Schrödinger equation also has disadvantages. Optimal control problems are typically defined within a finite simulation volume; for this reason, the application of absorbing boundary conditions (ABC) [Antoine et al., 2004, 2008] for the Bellman wavefunction is desirable. Because quantum mechanics allows for tunneling through barriers, the solutions obtained are likely to permit part of a Bellman wavefunction to pass through boundaries that should be impenetrable, raising questions about errors on the probability distribution.

2 | A challenging problem for optimal control: Zermelo's navigation problem

Zermelo's navigation problem [Zermelo, 1931] has become a classical problem for the field of optimal control. The original statement of the problem is

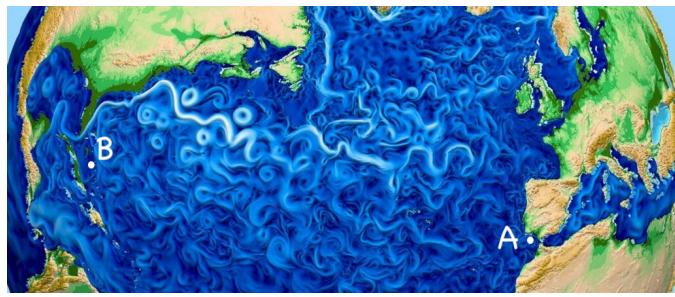


Figure 2.1: A visualization from the open-source E3SM code (image credit Argonne Leadership Computing Facility Science News) of an example setting for Zermelo's problem. A starting point (point A) and an ending point (point B) are marked. Optimal navigation between point A at Palos de la Frontera and point B San Salvador would identify an optimal route for a ship moving at a given speed on a voyage across the Atlantic ocean like that of Columbus' in 1492.

In einer unbegrenzten Ebene, in welcher die Windverteilung durch ein Vektorfeld als Funktion von Ort und Zeit gegeben ist, bewegt sich ein Fahrzeug mit konstanter Eigengeschwindigkeit relativ zur umgebenden Luftmasse. Wie muß das Fahrzeug gesteuert werden, um in kürzester Zeit von einem Ausgangspunkte zu einem gegebenen Ziel zu gelangen?

Typically the setting is described as a ship that is required to traverse between two points on the two-dimensional surface of a body of water in the presence of some current or wind (see Figure 2.1). In this work, we will generally refer to the current or wind as a background field. The optimal control requirement is that the ship starting at point A should reach the prescribed final point (point B) in the least possible time; Zermelo's problem has sometimes been reformulated to calculate an optimal energy path. Zermelo's assumption was that the ship would be capable of a fixed speed, so that only the angle is relevant in navigating an optimal path. A direct application of Zermelo's problem is found in the use of active drifters as ocean sensors [Ouimet and Cortés, 2014, Molchanov et al., 2015, Wiggert et al., 2022]. More complex variations of Zermelo's problem involve navigation in three-dimensional environments for active swimmers [Gunnarson et al., 2021, Nasiri and Liebchen, 2022, Nasiri et al., 2024], migrating animals [Hays et al., 2014], or flight [Techy, 2011].

When the original statement of Zermelo's problem is modified so that background field is taken to be constant in time, it is tractable for a wide range of solution types, including machine learning techniques. However there is no general approach that solves this problem with an evolving background field. Thus the question is raised: how significantly does the background field need to vary for machine learning solutions to be intractable. Zermelo's problem inherently has two different time-scales built in: one time-scale is associated with the evolution of the background field τ_{field} , while the other is associated with the ship τ_{ship} . In the limit where the ship moves fast

$$\tau_{\text{ship}} \ll \tau_{\text{field}}, \quad (2.1)$$

so that the background field can be assumed to be static. In the limit where the background field evolves rapidly

$$\tau_{\text{field}} \ll \tau_{\text{ship}}, \quad (2.2)$$

so that the background field can be assumed to be essentially stochastic, changing and then changing again before the ship can traverse further in a meaningful way. The difficult regime is thus the regime where both ship and background field evolve on the same time-scale

$$\tau_{\text{field}} \sim \tau_{\text{ship}}. \quad (2.3)$$

In this case an optimal path would be determined by the ship, and begun to be traversed. Mid-traversal, the background field would evolve so that the original path would no longer be optimal. A feedback loop exists between the ship and the non-linear PDEs from fluid dynamics that calculate the evolution of the background field, whether they are used to calculate a current or a wind.

A recent set of papers [Biferale et al., 2019, Buzzicotti et al., 2020] has addressed this kind of problem using turbulence – a setting where significant changes in the background field are implied – and reinforcement

learning methods. They also explore a variation on Zermelo's problem where the point B is taken to be another ship or a drifter, and so is allowed to move in time [Calascibetta et al., 2023a,b]. Biferale et al. [2019] addresses ship velocities from 20% of the maximal flow velocity up to the maximal flow velocity. Assuming that the background field changes on a characteristic length scale L , which the ship also has to traverse, we find an ordering

$$0.2v_{\text{field}} \leq v_{\text{ship}} \leq v_{\text{field}} , \quad (2.4)$$

$$0.2 \frac{L}{\tau_{\text{field}}} \leq \frac{L}{\tau_{\text{ship}}} < \frac{L}{\tau_{\text{field}}} . \quad (2.5)$$

This translates to time scales for the ship in the range

$$\tau_{\text{field}} < \tau_{\text{ship}} \leq 5\tau_{\text{field}} . \quad (2.6)$$

At the point where $\tau_{\text{ship}} = 5\tau_{\text{field}} \gg \tau_{\text{field}}$, the background field is approaching the stochastic situation that we mention above. We note that it would be interesting to examine a more symmetric time range, in which the velocity of the ship could also be slightly faster than the maximal flow velocity.

This collection of works [Biferale et al., 2019, Buzzicotti et al., 2020, Calascibetta et al., 2023a,b] demonstrates two things: (1) that reinforcement learning based on the actor-critic algorithm can be used to solve this type of problem, i.e. the learning process produces robust and stable solutions that result in the ship moving from position A to position B, and (2) using this machine learning method is much more computationally expensive for the case of an evolving background field than for a static background field. Methods that offer an alternative to traditional machine learning approaches, particularly methods with the potential to be more computationally efficient, are therefore desirable for the setting of background fields that evolve significantly.

3 | A Madelung Transform for the HJB equation

A variable transform exists that allows for the complex-valued Schrödinger equation to be written in the form of the Euler equations; this is called the Madelung Transform [Kambe, 2007, Zylberman et al., 2022, Carles et al., 2012, Gay-Balmaz and Tronci, 2020, Lü, 2015]. Here we recall the details of that transform, and interpret them in terms of the original HJB equation for optimal control. We begin with the basic Schrödinger-type equation:

$$i\hbar \frac{\partial}{\partial t} \psi(r, t) + \frac{\hbar^2}{2m} \nabla^2 \psi(r, t) - V(r) \psi(r, t) = 0 . \quad (3.1)$$

To divide this equation into real and imaginary parts, without loss of generality we express the wavefunction in terms of a phase ϕ and an amplitude A

$$\psi(r, t) = A(r, t) \exp(i\phi(r, t)) . \quad (3.2)$$

We then straightforwardly apply the chain-rule to express the time derivative of the wavefunction

$$\frac{\partial}{\partial t} \psi = \exp(i\phi) \frac{\partial}{\partial t} A + A \exp(i\phi) \frac{\partial}{\partial t}(i\phi) , \quad (3.3)$$

Dividing by the wavefunction, we find

$$\frac{1}{\psi} \frac{\partial}{\partial t} \psi = \frac{1}{A} \frac{\partial A}{\partial t} + i \frac{\partial \phi}{\partial t} . \quad (3.4)$$

Similarly, we use the amplitude and phase to express the spatial gradient of the wavefunction

$$\nabla \psi = \exp(i\phi) \nabla A + A \exp(i\phi) \nabla(i\phi) , \quad (3.5)$$

and then the Laplacian

$$\nabla^2 \psi = \nabla \cdot \nabla \psi = \exp(i\phi) \nabla^2 A + A \exp(i\phi) \nabla^2(i\phi) \quad (3.6)$$

$$+ A(\exp(i\phi) \nabla(i\phi)) \cdot (\nabla(i\phi)) + 2 \exp(i\phi) (\nabla A) \cdot (\nabla(i\phi)) . \quad (3.7)$$

Dividing by the wavefunction, we find

$$\frac{1}{\psi} \nabla^2 \psi = \left[\frac{1}{A} \nabla^2 A - (\nabla \phi)^2 \right] + i \left[\nabla^2 \phi + \frac{2}{A} (\nabla A) \cdot (\nabla \phi) \right]. \quad (3.8)$$

Using these in the Schrödinger equation (3.1) and also dividing by ψ , we obtain

$$\frac{i\hbar}{A} \frac{\partial A}{\partial t} - \hbar \frac{\partial \phi}{\partial t} + \frac{\hbar^2}{2m} \left[\frac{1}{A} \nabla^2 A - (\nabla \phi)^2 \right] + \frac{i\hbar^2}{2m} \left[\nabla^2 \phi + \frac{2}{A} (\nabla A) \cdot (\nabla \phi) \right] - V(r) = 0. \quad (3.9)$$

Separating this into the real and imaginary parts we obtain two coupled equations:

$$\hbar \frac{\partial}{\partial t} \phi - \frac{\hbar^2}{2mA} \nabla^2 A + \frac{\hbar^2}{2m} (\nabla \phi)^2 + V(r) = 0, \quad (3.10)$$

$$\frac{i\hbar}{A} \frac{\partial A}{\partial t} + \frac{i\hbar^2}{2m} \left[\nabla^2 \phi + \frac{2}{A} (\nabla A) \cdot (\nabla \phi) \right] = 0. \quad (3.11)$$

We multiply the equation for the amplitude by $2mA^2/i\hbar$ and apply the chain rule to the derivatives of A .

$$2mA \frac{\partial A}{\partial t} + \hbar A^2 \nabla^2 \phi + 2\hbar A (\nabla A) \cdot (\nabla \phi) = 0. \quad (3.12)$$

$$\frac{\partial}{\partial t} mA^2 + \hbar A^2 \nabla^2 \phi + \hbar \nabla A^2 \cdot \nabla \phi = 0. \quad (3.13)$$

We then identify a mass probability density

$$\rho_B = m\psi\psi^* = mA^2, \quad (3.14)$$

and a mass flux

$$\mathbf{j} = (\hbar/2i)(\psi^* \nabla \psi - \psi \nabla \psi^*) = \hbar A^2 \nabla \phi. \quad (3.15)$$

A velocity field can be recovered from these two quantities

$$\mathbf{u}_B = \mathbf{j}/\rho_B = \nabla \left(\frac{\hbar}{m} \phi \right) = \nabla \Phi. \quad (3.16)$$

Here in the final equality we have absorbed the constants \hbar/m into a renormalized phase Φ . The equation for the evolution of A (3.13) becomes

$$\frac{\partial}{\partial t} mA^2 + (mA^2) \nabla^2 \left(\frac{\hbar}{m} \phi \right) + \nabla mA^2 \cdot \nabla \frac{\hbar}{m} \phi = 0. \quad (3.17)$$

$$\frac{\partial}{\partial t} \rho_B + \nabla \rho_B \cdot \mathbf{u}_B + \rho_B \nabla \cdot \mathbf{u}_B = 0. \quad (3.18)$$

$$\frac{\partial}{\partial t} \rho_B + \nabla \cdot (\rho_B \mathbf{u}_B) = 0. \quad (3.19)$$

This final equation (3.19) is the usual continuity equation for a compressible fluid.

We rearrange eq. (3.10) into an equation for the velocity. We first divide by m to obtain an equation for the renormalized phase Φ :

$$\frac{\hbar}{m} \frac{\partial}{\partial t} \phi - \frac{\hbar^2}{2m^2 A} \nabla^2 A + \frac{\hbar^2}{2m^2} (\nabla \phi)^2 = -\frac{1}{m} V(r). \quad (3.20)$$

We then take the gradient

$$\frac{\partial}{\partial t} \nabla \frac{\hbar}{m} \phi - \nabla \frac{\hbar^2}{2m^2 A} \nabla^2 A + \nabla \frac{\hbar^2}{2m^2} (\nabla \phi)^2 = -\nabla \frac{1}{m} V(r), \quad (3.21)$$

$$\frac{\partial}{\partial t} \nabla \Phi + (\nabla \Phi) \cdot \nabla (\nabla \Phi) = -\nabla \left[\frac{1}{m} V(r) + \frac{\hbar^2}{2m^2 A} \nabla^2 A \right]. \quad (3.22)$$

Changing the notation to that of the velocity, we find

$$\frac{\partial}{\partial t} \mathbf{u}_B + \mathbf{u}_B \cdot \nabla \mathbf{u}_B = -\nabla \mu. \quad (3.23)$$

In the final form of this equation, we have defined a new scalar potential $\mu = V(r)/m - (\hbar^2/2m^2)\nabla^2 A/A$. Eq. (3.23) is the nonlinear fluid equation for the evolution of velocity. Taken together, eqs. (3.19) and (3.23) are an instance of the Euler equations in fluid dynamics. The quantity μ encapsulates all of the optimization that the Bellman potential did in the Schrödinger formulation – or that the Bellman value function originally contained. The form of this scalar potential is the same as a pressure term, a gravitational potential, or a chemical potential. We will refer to μ as the *Bellman pressure*. Using the Bellman pressure, the Bellman value function can now be interpreted as a pressure force on the flow.

The formulation of the HJB equation as a fluid governed by the Euler equations provides a new perspective on the field of fluid control. The Bellman density ρ_B provides a mass-weighted probability for the Bellman wavefunction ψ . Conservation of probability in the Schrödinger formulation of the optimal control problem thus translates to conservation of mass in the Euler fluid. The Bellman velocity field \mathbf{u}_B defines the local movement of the probability density in the Schrödinger formulation; making this velocity explicit implies that it can be varied in either the Schrödinger formulation or the Euler formulation.

4 | Formulation of Zermelo's problem as multi-fluid equations

Applying the Euler formulation in Section 3 to Zermelo's problem, we have two coupled equations that describe the optimal path of the ship

$$\frac{\partial}{\partial t}\rho_B + \nabla \cdot (\rho_B \mathbf{u}_B) = 0, \quad (4.1)$$

$$\frac{\partial}{\partial t}\mathbf{u}_B + \mathbf{u}_B \cdot \nabla \mathbf{u}_B = -\nabla \mu. \quad (4.2)$$

These equations are coupled to the incompressible Euler equations for the background field

$$\nabla \cdot \mathbf{u}_0 = 0, \quad (4.3)$$

$$\frac{\partial}{\partial t}\mathbf{u}_0 + \mathbf{u}_0 \cdot \nabla \mathbf{u}_0 = -\nabla p. \quad (4.4)$$

Here we have adopted the null subscript to denote the air or fluid flow at the surface of a body of water. Using an incompressible fluid seems like a reasonable first step to calculate the background field evolution, but compressible flows, 3D flows, and flows with additional force terms could also be examined.

For the problem as formulated in eqs. (4.1) – (4.4), the use of common numerical methods for CFD could be brought to bear, as a single online calculation. Namely the equations for the background field eqs. (4.3) – (4.4) could be evolved one time step. Then the Bellman pressure μ could be calculated based on the solution for the background field. Finally, the equations for the ship's path eqs. (4.1) – (4.2) could be evolved a time step by the solver. The ship can thus be viewed as a *passive* second fluid – more complex than a passive scalar, but not significantly so. As a Bellman fluid, the ship would initially be concentrated at a given position; it would then be transported and undergo diffusion using contributions to μ from the actual fluid. The Bellman pressure would need to be solved for separately from the fluid equations, as is commonly done, e.g. for gravitational accelerations with a Poisson solver. The practicality of this multi-fluid approach to optimal control depends on the efficiency of that solution, which could itself be computationally demanding.

5 | Broader implications for the field of fluid control

To approach the solution for the Bellman pressure, here we briefly note that this concept is parallel to ideas from the field of fluid animation. Notable achievements in fluid animation include examples like the fire from the dragon's mouth in *Harry Potter and the Goblet of Fire* (Warner Bros. 2005), the ocean simulations in *Moana* (Disney 2016) [Byun and Stomakhin, 2017], and fire, explosions and water effects in *Avatar: The Way of Water* (20th Century 2022) [Edholm et al., 2023, Gowanlock, 2012, 2021]. Computer graphics for fluids have to look physically plausible to the human eye, i.e. smaller length scales need to be generated by the relevant physical equations of fluid dynamics. However, the smallest scales, smaller than the level noticeable by the human eye, do not need to be included in fluid animation; this is similar to the philosophy of large eddy simulations (LES) [Schmidt, 2015, Lesieur and Metais, 1996] where length scales

smaller than some chosen cut-off are filtered out of the equations, and their effect on the large scales is modeled using a sub-grid scale model.

The existence of successful fluid animations thus demonstrates that solution of the Bellman pressure should be feasible. Most fluid animation techniques can be classified into one of two general types:¹ (1) proportional-derivative controllers that guide the fluid body using additional “ghost force” terms that are designed based on a distance measure between the current fluid shape and the desired one [e.g. [Feng and Liu, 2018](#), [Thürey et al., 2009](#)], (2) optimal controllers that formulate the problem as a space-time optimization over the space of possible control forces constrained by the fluid governing equations [e.g. [Treuille et al., 2003](#), [Pan and Manocha, 2017](#)]. These methods operate using a concept of “keyframes”, i.e. a sequence of states where the largest scales of the desired flow structure, e.g. the outline of the dragon’s breath, is drawn into frames of the movie. The solution from each frame to the next thus solves a problem like Zermelo’s problem, where the flow has to move between a point A and a point B. The optimization of the Bellman pressure μ required by our method echoes the fluid animation concepts of calculating external forces that move the fluid to match a desired shape at specific times. The details of fluid animation methods may point a way forward in the development of an algorithm for Zermelo’s problem as well.

6 | Summary

The transformation of the HJB equation into a Schrödinger equation allows for the application of general scientific computing methods to be applied to otherwise intractable optimal control problems. In this white paper, we outline a further transformation to the Euler equations which could be convenient, especially for optimal navigation problems which could then be expressed as a multi-fluid problem. In examining these transformations, we have drawn an analogy between the Bellman value function and a pressure-like term that we call the Bellman pressure. We have pointed out that optimal fluid controllers used for fluid animation are conceptually equivalent to solving for a Bellman pressure. This suggests that such a solver would be numerically feasible. Using a multi-fluid set of equations, we could solve for the evolving fluid dynamics, as well as the coupled evolving optimal path of the ship using the same numerical infrastructure. This would allow, for example, a framework for optimal path calculation to be incorporated in a code like Energy Exascale Earth System Model (E3SM) [[Rasch et al., 2019](#)] to determine optimal paths through, e.g. evolving weather patterns across the globe.

Acknowledgements

This material is based upon work supported by the National Science Foundation under grant no. PHY-1907876. Part of this work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-TR-2010440.

7 | References

Peter A Forsyth and George Labahn. Numerical methods for controlled hamilton-jacobi-bellman pdes in finance. *Journal of Computational Finance*, 11(2):1, 2007.

Hilbert J Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of statistical mechanics: theory and experiment*, 2005(11):P11011, 2005.

Franziska Kühn. Viscosity solutions to Hamilton–Jacobi–Bellman equations associated with sublinear Lévy (-type) processes. *ALEA*, 16:531–559, 2019.

Martino Bardi, Italo Capuzzo Dolcetta, et al. *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations*, volume 12. Springer, 1997.

Luis G Crespo and Jian-Qiao Sun. Stochastic optimal control via Bellman’s principle. *Automatica*, 39(12):2109–2114, 2003.

¹Deep/reinforcement learning methods have also been applied [[Brunton et al., 2020](#)].

Marianne Akian, Jan Blechschmidt, Nikolai D Botkin, Max Jensen, Axel Kröner, Athena Picarelli, Iain Smears, Karsten Urban, Mickaël D Chekroun, Roland Herzog, et al. *Hamilton-Jacobi-Bellman equations: numerical methods and applications in optimal control*, volume 21. Walter de Gruyter GmbH & Co KG, 2018.

Mario Annunziato, Alfio Borzì, Fabio Nobile, and Raul Tempone. On the connection between the Hamilton-Jacobi-Bellman and the Fokker-Planck control frameworks. *Applied Mathematics*, 2014.

Michael G Crandall. Viscosity solutions: a primer. In *Viscosity Solutions and Applications: Lectures given at the 2nd Session of the Centro Internazionale Matematico Estivo (CIME) held in Montecatini Terme, Italy, June 12–20, 1995*, pages 1–43. Springer, 2006.

Michael D Schneider, Caleb Miller, George F Chapline, Jane Pratt, and Dan Merl. Optimal control from inverse scattering via single-sided focusing. *arXiv preprint arXiv:2212.00249*, 2022.

Alexis MH Teter, Abhishek Halder, Michael D Schneider, Alexx S Perloff, Jane Pratt, Conor M Artman, and Maria Demireva. Control-affine Schrödinger Bridge and Generalized Bohm potential. *arXiv preprint arXiv:2508.08511*, 2025.

Xavier Antoine, Christophe Besse, and Vincent Mouysset. Numerical schemes for the simulation of the two-dimensional Schrödinger equation using non-reflecting boundary conditions. *Mathematics of computation*, 73(248):1779–1799, 2004.

Xavier Antoine, Anton Arnold, Christophe Besse, Matthias Ehrhardt, and Achim Schädle. A review of transparent and artificial boundary conditions techniques for linear and nonlinear Schrödinger equations. 2008.

Ernst Zermelo. Über das Navigationsproblem bei ruhender oder veränderlicher Windverteilung. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, 11(2):114–124, 1931.

Michael Ouimet and Jorge Cortés. Coordinated rendezvous of underwater drifters in ocean internal waves. In *53rd IEEE Conference on Decision and Control*, pages 6099–6104. IEEE, 2014.

Artem Molchanov, Andreas Breitenmoser, and Gaurav S Sukhatme. Active drifters: Towards a practical multi-robot system for ocean monitoring. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 545–552. IEEE, 2015.

Marius Wiggert, Manan Doshi, Pierre FJ Lermusiaux, and Claire J Tomlin. Navigating underactuated agents by hitchhiking forecast flows. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 2417–2424. IEEE, 2022.

Peter Gunnarson, Ioannis Mandralis, Guido Novati, Petros Koumoutsakos, and John O Dabiri. Learning efficient navigation in vortical flow fields. *Nature communications*, 12(1):7143, 2021.

Mahdi Nasiri and Benno Liebchen. Reinforcement learning of optimal active particle navigation. *New Journal of Physics*, 24(7):073042, 2022.

Mahdi Nasiri, Edwin Loran, and Benno Liebchen. Smart active particles learn and transcend bacterial foraging strategies. *Proceedings of the National Academy of Sciences*, 121(15):e2317618121, 2024.

Graeme C Hays, Asbjørn Christensen, Sabrina Fossette, Gail Schofield, Julian Talbot, and Patrizio Mariani. Route optimisation and solving Zermelo’s navigation problem during long distance migration in cross flows. *Ecology letters*, 17(2):137–143, 2014.

Laszlo Techy. Optimal navigation in planar time-varying flow: Zermelo’s problem revisited. *Intelligent Service Robotics*, 4(4):271–283, 2011.

Luca Biferale, Fabio Bonaccorso, Michele Buzzicotti, Patricio Clark Di Leoni, and Kristian Gustavsson. Zermelo’s problem: optimal point-to-point navigation in 2d turbulent flows using reinforcement learning. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(10), 2019.

Michele Buzzicotti, Luca Biferale, Fabio Bonaccorso, Patricio Clark Di Leoni, and Kristian Gustavsson. Optimal control of point-to-point navigation in turbulent flows using reinforcement learning. In *AIRO@AI*IA*, pages 39–43, 2020.

Chiara Calascibetta, Luca Biferale, Francesco Borra, Antonio Celani, and Massimo Cencini. Optimal tracking strategies in a turbulent flow. *Communications Physics*, 6(1):256, 2023a.

Chiara Calascibetta, Luca Biferale, Francesco Borra, Antonio Celani, and Massimo Cencini. Taming Lagrangian chaos with multi-objective reinforcement learning. *The European Physical Journal E*, 46(3):9, 2023b.

Tsutomu Kambe. *Elementary fluid mechanics*. World Scientific, 2007.

Julien Zylberman, Giuseppe Di Molfetta, Marc Brachet, Nuno F Loureiro, and Fabrice Debbasch. Quantum simulations of hydrodynamics via the Madelung transformation. *Physical Review A*, 106(3):032408, 2022.

Rémi Carles, Raphaël Danchin, and Jean-Claude Saut. Madelung, Gross–Pitaevskii and Korteweg. *Nonlinearity*, 25(10):2843, 2012.

François Gay-Balmaz and Cesare Tronci. Madelung transform and probability densities in hybrid quantum–classical dynamics. *Nonlinearity*, 33(10):5383, 2020.

Xing Lü. Madelung fluid description on a generalized mixed nonlinear schrödinger equation. *Nonlinear Dynamics*, 81(1):239–247, 2015.

Dong Joo Byun and Alexey Stomakhin. Moana: crashing waves. In *ACM SIGGRAPH 2017 Talks*, pages 1–2. 2017.

John Edholm, Alexey Stomakhin, Rahul Deshpabhu, David Caeiro Cebrian, Florian Hu, and Caitlin Pope. Fire and explosions in Avatar: The Way of Water. In *ACM SIGGRAPH 2023 Talks*, pages 1–2. 2023.

Jordan Gowanlock. *Particles and waves: computational fluid dynamics and film theory*. PhD thesis, Carleton University, 2012.

Jordan Gowanlock. *Animating unpredictable effects: Nonlinearity in Hollywood’s R&D complex*. Springer Nature, 2021.

Wolfram Schmidt. Large eddy simulations in astrophysics. *Living Reviews in Computational Astrophysics*, 1(1):2, 2015.

Marcel Lesieur and Olivier Metais. New trends in Large-Eddy Simulations of turbulence. *Annual review of fluid mechanics*, 28(1):45–82, 1996.

Steven L Brunton, Bernd R Noack, and Petros Koumoutsakos. Machine learning for fluid mechanics. *Annual review of fluid mechanics*, 52(1):477–508, 2020.

Gang Feng and Shiguang Liu. Detail-preserving sph fluid control with deformation constraints. *Computer Animation and Virtual Worlds*, 29(1):e1781, 2018.

Nils Thürey, Richard Keiser, Mark Pauly, and Ulrich Rüde. Detail-preserving fluid control. *Graphical Models*, 71(6):221–228, 2009.

Adrien Treuille, Antoine McNamara, Zoran Popović, and Jos Stam. Keyframe control of smoke simulations. In *ACM SIGGRAPH 2003 Papers*, pages 716–723. 2003.

Zherong Pan and Dinesh Manocha. Efficient solver for spacetime control of smoke. *ACM Transactions on Graphics (ToG)*, 36(4):1, 2017.

Philip J Rasch, S Xie, P-L Ma, W Lin, H Wang, Q Tang, SM Burrows, P Caldwell, K Zhang, RC Easter, et al. An overview of the atmospheric component of the energy exascale earth system model. *Journal of Advances in Modeling Earth Systems*, 11(8):2377–2411, 2019.