Learning high-fidelity GW models from numerical relativity data

Scott Field

Department of Mathematics, U. Mass Dartmouth

ICERM Workshop
Nov 18, 2020

Overview

Part I (25 minutes): Overview, usage, models

Part II (20 minutes): (tutorial) Methods for building a 1-dimensional model

Part III (25 minutes): (tutorial) Building a 1-dimensional model

Part IV (45 minutes): (tutorial) gwsurrogate, SurfinBH, and binaryBHexp (Vijay Varma)

Advances and Challenges in Computational General Relativity

MAIN PAGE
HOTELS AND TRAVEL
SCHEDULE
PARTICIPANTS
REGISTRATION
POSTER SESSION
FINANCIAL SUPPORT

Division of Applied Mathematics
Brown University
182 George Street
Providence, RI 02912
Telephone: (401) 863-2115
Fax: (401) 863-1355
dam@dam.brown.edu

Advances and Challenges in Computational General Relativity

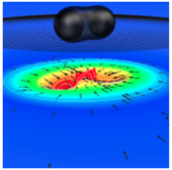

SCIENTIFIC PROGRAM

All talks will take place in Barus & Holley, room 190.

For a PDF version of the program click here

Friday, May 20, 2011

Time	Speaker	Affiliation	Title
08:00 - 08:45	Registration and light breakfast		Outside rosom 190, Barus & Holley
08:45 - 09:00	Opening remarks		

MAIN PAGE
HOTELS AND TRAVEL
SCHEDULE
PARTICIPANTS
REGISTRATION
POSTER SESSION
FINANCIAL SUPPORT

Division of Applied Mathematics
Brown University
182 George Street
Providence, RI 02912
Telephone: (401) 863-2115
Fax: (401) 863-1355
dam@dam.brown.edu

Advances and Challenges in Computational General Relativity

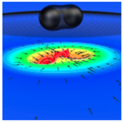

SCIENTIFIC PROGRAM

All talks will take place in Barus & Holley, room 190.

For a PDF version of the program click here

Friday, May 20, 2011

Time	Speaker	Affiliation	Title
08:00 - 08:45	Registration and light breakfast		Outside rosom 190, Barus & Holley
08:45 - 09:00	Opening remarks		
09:00 - 09:30	Peter Diener	Louisiana State U	The Effective Source Approach to the Self-force Problem (PDF)
09:30 - 10:00	Scott Field	Brown U	A Discontinuous Galerkin Method for BSSN-Type Systems (PDF)
10:00 - 10:30	Manuel Tiglio	U of Maryland	Reduced Basis in General Relativity: Select, Solve, Represent, Predict (PDF)

Collaborators

Surrogate modeling methods have been developed and refined by many people over since 2011...

Jonathan Blackman, Chad Galley, Vijay Varma, Nur Rifat, Gaurav Khanna, Frank Herrmann, Jan Hesthaven, Evan Ochsner, Manuel Tiglio, Harbir Antil, Ricardo Nochetto, Jason Kaye, Bela Szilagyi, Mark Scheel, Dan Hemberger, Rory Smith, Kent Blackburn, Carl Haster, Michael Purrer, Stephen Lau, Saul Teukolsky, Vivien Raymond, Patricia Schmidt, Mike Boyle, Larry Kidder, Harald Pfeiffer, Davide Gerosa, Leo Stein, Tousif Islam, Feroz Shaik

blue = significant contributors to gwsurrogate code

...and *many* simulations

SIMULATING EXTREME SPACETIMES

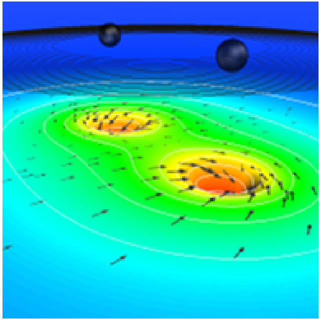

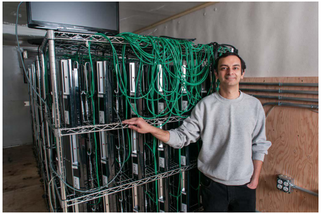

Outline

Motivation/Overview

Gravitational waveform generation from compact binary coalescences is a computational bottleneck for...

- Template-based detection algorithms
- Parameter estimation
- Calibration of phenomenological or effective merger models (with NR)

Motivation/Overview

Gravitational waveform generation from compact binary coalescences is a computational bottleneck for...

- Template-based detection algorithms
- Parameter estimation
- Calibration of phenomenological or effective merger models (with NR)

Strategy for parameterized waveform models

- Train an accurate and fast-to-evaluate surrogate model
- The model is built entirely from simulation data
- Only possible given the recent progress made in numerical relativity
- *NOT* reduced physics
    - Surrogate converges to underlying model (NR) with more waveform data
    - Trade-off: model only valid in its training (temporal/parametric) interval

Other approaches to speedup

Computational bottlenecks due to waveform generation costs are ubiquitous. Alternative solutions include...

- Closed-form & phenomenological models (Phenom{A,B,C,D,P,Pv2}, effective-one-body)
- Algorithmic and hardware optimization of pipelines (e.g. GstLAL, PyCBC)
- Extensive, model-specific optimizations (e.g. Devine, Etienne, McWilliams)
- GPU acceleration (see tutorial by Michael Katz)
- NR-based parameter estimation (see talk by Richard O'Shaughnessy)
- And more!

What is a surrogate model?

Surrogate(Merriam-Webster): one that serves as a substitute – mimics behavior of the full, underlying model for a fixed range of the parameter and physical variables

What is a surrogate model?

Surrogate(Merriam-Webster): one that serves as a substitute – mimics behavior of the full, underlying model for a fixed range of the parameter and physical variables

Features

- Surrogate will converge to underlying model with more training data
- Only reproduces outputs of interest (waveforms, remnant values, etc)
- Should be viewed as a waveform acceleration technique

Decisions

- At which parameters should one evaluate the underlying model?
- How to tie together these samples?
- Often times different methods (e.g. SVD vs greedy; fits vs GPR) will result in similar surrogate model quality – choices may just be a matter of familiarity or convenience.

Examples

- Machine learning, fits/interpolation, reduced order modeling
- At least for this talk, ROM = surrogate

Why do we need surrogate models?

- They are nearly indistinguishable from the underlying model
- EOB surrogate models enable speed up factors of between 10^2 - 10^3
- NR surrogates speedups $\approx 10^7$ (0.01 seconds vs $\approx$ 1 week)
- Due to these speedups, surrogates enable new kinds of studies to be carried out
  - Typical Bayesian inference run requires $> 10^6$ model evaluations

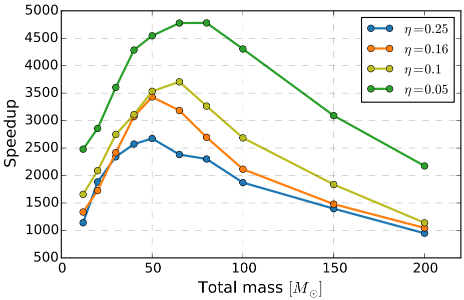

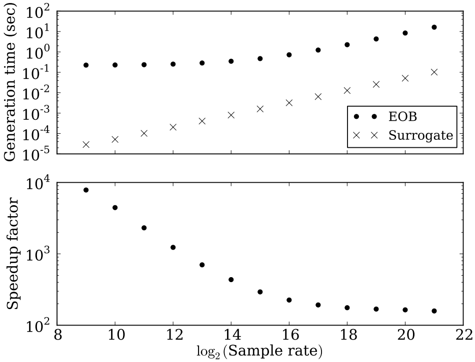

Surrogate models (without matter)

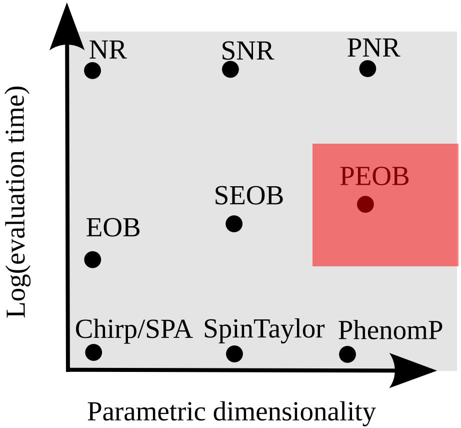

- Closed-form waveform models
  - Cannon et al. (2010, 2012, 2013), Field et al. (2011, 2012), Kaye (2012), Smith et al. (2013, 2016), Doctor et al. (2017), Chua et al. (2018)
- (Multi-mode) Effective one body (EOB)
  - Field et al., (2014), Purrer (2014, 2016), Lackey et al. (2019), Cotesta et al. (2020)
- Multi-mode numerical relativity
  - Blackman et al., 2015 (non-spinning),
  - Blackman et al., 2017 (5d subspace),
  - Blackman et al., 2017 (full 7d , $q \leq 2$)
  - Varma et al., 2019 (enlarged 7d, $q \leq 4$)
  - Varma et al., 2019 (Hybridized, aligned spin)

Tidal models and $q \leq 10^4$ have also been built

Surrogates in LIGO-Virgo data analysis

- Accelerate waveform generation by factors of 10^2 (EOB models build by Purrer and Cotesta; described by ODEs) to 10^8 (NR models; described by PDEs)
- EOB ROMs are extensively used as part of the LSC's parameter-estimation efforts as well as template-bank detection
- NR surrogates have been used in for specific BBH events
- Surrogate models have been essential to the widespread use of both EOB and NR waveforms for realistic data analysis efforts with LSC data

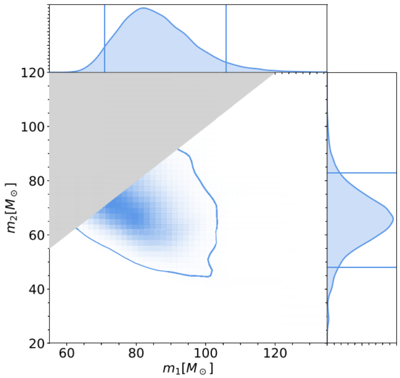

FIG. 2. Posterior distributions for the progenitor masses of GW190521 according to the NRSur7dq4 waveform model. The 90% credible regions are indicated by the solid contour in the joint distribution and by solid vertical and horizontal lines in the marginalized distributions.

Who's using surrogates? (partial list)

- Studies of gravitational wave memory (Lasky et al; PRL 2020)
- Training neural networks (Wei et al; Physics Letters B 2020)
- Validating searches for primordial BHs (Nitz et al; arXiv:2007.03583)
- Measuring kicks (Varma et al; PRL 2020)
- Building/assessing other models (Garca-Quirs et al; PRD 2020)
- Studying systematics of subdominant modes (Shaik et al; PRD 2020)
- Analyzing GW190412 (Islam et al; arXiv:2010.04848)

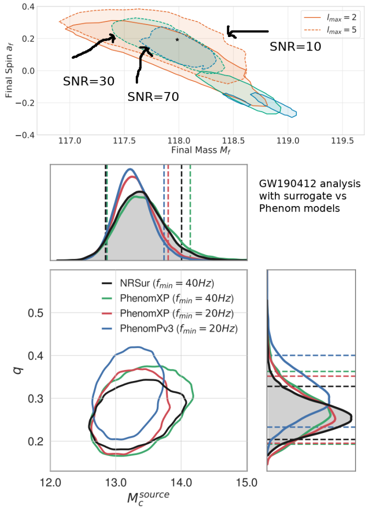

Surrogates are great! What models can I use?

See Vijay Varma's tutorial next for a full introduction to models for the waveform, dynamics, and remnant properties

Outline

GWSurrogate Python package

Goals:

- Surrogate-building codes and data are model-specific (sometimes very different)
- GWSurrogate: easy to install, easy to use, Python{2,3}-based
- Current catalog of surrogate models + data access tools

Why not just use LALSimulation?

GWSurrogate Python package

Goals:

- Surrogate-building codes and data are model-specific (sometimes very different)
- GWSurrogate: easy to install, easy to use, Python{2,3}-based
- Current catalog of surrogate models + data access tools

Why not just use LALSimulation? Some models will be ported, but...

- Not everyone can or should need to install LALSimulation to use surrogates
- Its unlikely that for each new surrogate there will be LALSimulation counterpart
- Having multiple codes to evaluate the same model is good for the community
- GWSurrogate API allows access of modes, basis functions, fits, and other surrogate data
- More than just waveforms! Dynamics, remnant properties (SurfinBH), etc...

GWSurrogate catalog

Installation:

```
>>> pip install gwsurrogate
```

GWSurrogate catalog

Installation:

```
>>> pip install gwsurrogate
```

Query the catalog:

```
>>> import gwsurrogate as gws
>>> gws.catalog.list(verbose=True)
NRSur7dq4
  url: https://zenodo.org/record/3348115/files/NRSur7dq4.h5
  md5 hash: 8e033ba4e4da1534b3738ae51549fb98
  Description: Surrogate model for precessing binary black holes with mass ratios q<=4
  and spin magnitudes <=0.8. This model is presented in Varma et al. 2019,
  arxiv:1905.09300. All ell<=4 modes are included. The spin and frame dynamics
  are also modeled.
  References: https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.1.033015
```

10 surrogate models are available. Each has a name, dataset url, description, and reference.

GWSurrogate catalog

```
>>> gws.catalog.list(verbose=False)
EOBNRv2
SpEC_q1_10_NoSpin
SpEC_q1_10_NoSpin_linear
SpEC_q1_10_NoSpin_linear_alt
NRSur4d2s_TDROM_grid12
NRSur4d2s_FDROM_grid12
NRHybSur3dq8
NRSur7dq4
NRHybSur3dq8Tidal
EMRISur1dq1e4
```

Lets look at some current and planned models...

- Name: SpEC_q1_10_NoSpin
- Nonspinning, $1 \leq q \leq 10$, $\ell \leq 8$, 22 NR training waveforms
- **Top**: Waveform differences between the two highest SpEC resolutions (black circles), the full surrogate and SpEC (red squares), and leave-one-out trial surrogates and SpEC (blue triangles).
- **Bottom**: The $(2,2)$ mode is shown for the largest surrogate error $q \approx 2$
- Not in LALSimulation (LVC code)

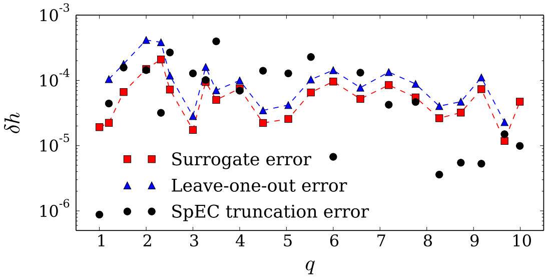

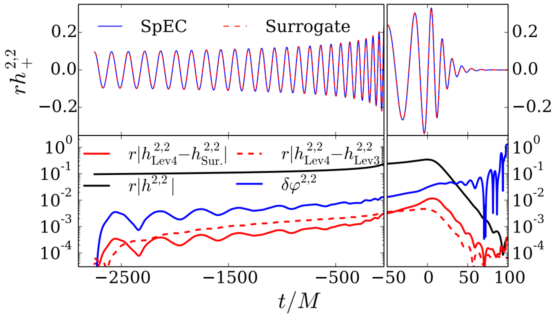

- Name: NRHybSur3dq8{Tidal}
- Hybridized spin-aligned model, $1 \leq q \leq 8$, $\ell \leq 4 + (5,5)$, spins < 0.8
- **Top**: 104 NR training waveforms sampling 3d space
- **Bottom**: Histogram of errors (last 4000M) for NR, NRHybSur3dq8, and SEOBNRv4HM
- Surrogate errors are cross-validation
- NRHybSur3dq8(-Tidal) is (is not) in LALSimulation
- GWSurrogate version of NRHybSur3dq8 has passed LVC code review

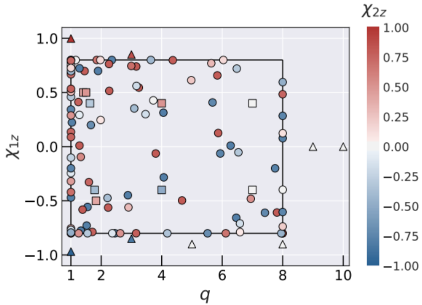

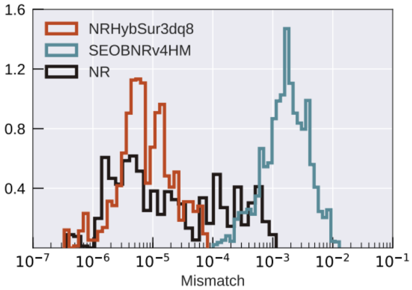

- Name: NRSur7dq4
- Generically precessing model, $1 \leq q \leq 4$, $\ell \leq 4$, spins < 0.8
- 1528 NR training waveforms sampling 7d space
- **Top**: (2,2) and (2,1) modes in inertial frame
- **Bottom**: Histogram of cross-validation errors
- NRSur7dq4 is in LALSimulation
- GWSurrogate version of has passed LVC code review, and includes dynamics

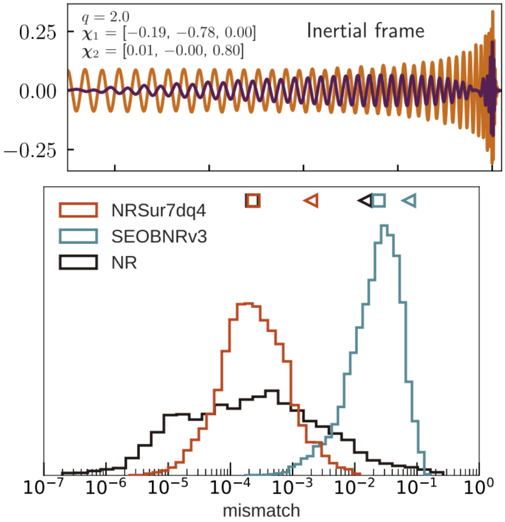

- Name: EMRISur1dq1e4
- Perturbation theory model, nonspinning, $3 \leq q \leq 10^4$, $\ell \leq 5$
- Adiabatic inspiral driven by GW radiation, a latestage geodesic plunge, Ori-Thorne transition trajectory between the two
- Training data generated with Gaurav Khanna's Teukolsky solver
- Le Tiec et al. (2011), Zimmerman et al. (2016), and others found certain NR and perturbation theory results agree surprisingly well at modest mass ratios. Waveforms seem to match too!
- EMRISur1dq1e4 is also in the BHPTK

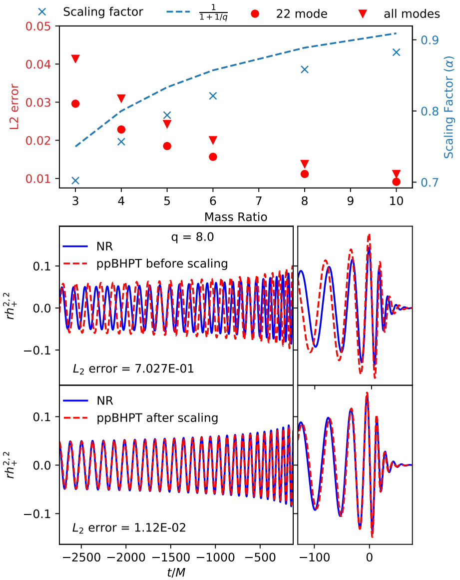

Eccentric model (Coming soon!)

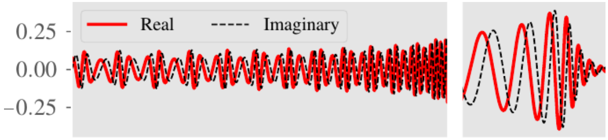

- Name: TBD
- Eccentric NR, nonspinning, $q \approx 1$, eccentricity non-zero at merger
- Tousif Islam, Vijay Varma, SF +
- **Top**: (2,2) mode in inertial frame
- **Bottom**: Cross-validation mismatch errors over parameter space

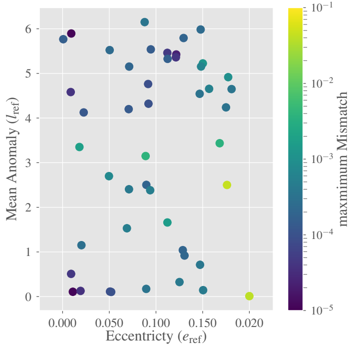

Remarks

- Surrogate and reduced order modeling offers an exciting new approach to overcome a variety of computationally challenging problems of GW physics
- Publicly available surrogate evaluation package GWSurrogate
  - Numerous Jupyter tutorial notebooks
  - Active development (version 1.0.7 released on Saturday)
  - Code hosted on github
  - Long-term plans: better documentation, more models, open/solve issues

Future outlook

- Extending the range of validity of NR & IMRI/EMRI surrogates
- As new models are built they will be included into the surrogate catalog
- Contributions are welcomed! If you've built a surrogate model, we can happily add it to the catalog

Surrogates are great! But none have been built for my favorite model. What should I do?

Outline

Lets build a 1D model

Tutorial location: `https://github.com/vijayvarma392/ICERM_workshop`

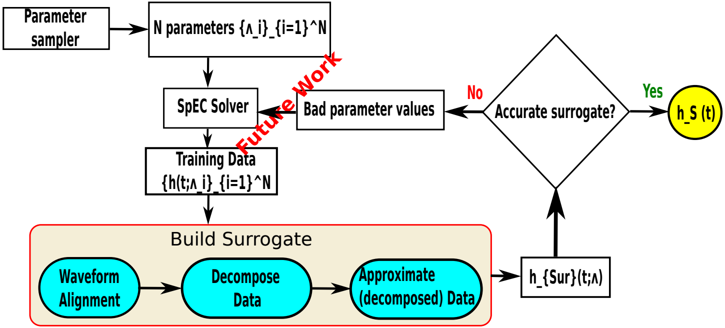

"top-level" view of surrogate model building.

The setup

Choose your favorite 1-dimensional gravitational-wave model:

$$\begin{aligned} h(t,\theta,\phi;q) &= h_+(t,\theta,\phi;q) - \mathrm{i}h_\times(t,\theta,\phi;q) \\ &= \sum_{\ell=2}^{\infty}\sum_{m=-\ell}^{\ell} h^{\ell m}(t;q)\,{}_{-2}Y_{\ell m}(\theta,\phi)\ , \end{aligned}$$

- θ and ϕ are angles for the direction of propagation away from the source.
- q is the mass ratio.
- ${}_{-2}Y_{\ell m}$ are the harmonic functions

We will build a model for $h^{\ell m}(t;q)$

Strategy for parameterized problems

Collect training data

- *Training set*: Evaluate the model at N values of q, giving us a N

$$\{h^{\ell m}(t; q_i)\}_{i=1}^{N}$$

  snapshots of the model
- *Training region*: $q \in [q_{\min}, q_{\max}]$, $t \in [t_{\min}, t_{\max}]$

Strategy for parameterized problems

Collect training data

- *Training set*: Evaluate the model at N values of q, giving us a N

$$\{h^{\ell m}(t;q_i)\}_{i=1}^{N}$$

  snapshots of the model
- *Training region*: $q \in [q_{\min}, q_{\max}]$, $t \in [t_{\min}, t_{\max}]$

Train the model

Train a surrogate model $h_{\mathrm{S}}^{\ell m}(t;q)$ such that $h^{\ell m}(t;q) \approx h_{\mathrm{S}}^{\ell m}(t;q)$ *within the training region*

This is a very general "learning from data" paradigm used in many fields of science and engineering

Read the fine print

Two key limitations of surrogate modeling...

Read the fine print

Two key limitations of surrogate modeling...

Drawback I: We must already have access to a model in order to build the surrogate.
Typical usage: the underlying model is too slow, the surrogate should be much faster to evaluate

Read the fine print

Two key limitations of surrogate modeling...

Drawback I: We must already have access to a model in order to build the surrogate.
Typical usage: the underlying model is too slow, the surrogate should be much faster to evaluate

Drawback II: We are only guaranteed the surrogate is accurate in the training region
Typical usage: Build the model for as large of a region as one expects to need

Methods for 1D (this session)

We will consider "traditional" methods. These are well-suited for 1D, 3D, and beyond. Some places they appeared include...

- Closed-form models: Cannon et al. (2010, 2012, 2013), Field et al. (2011, 2012), Kaye (2012), Smith et al. (2013, 2016)
- Nonspinning, multimode effective one body: Field et al., (PRX 2014)
- Spinning EOB: Purrer, (CQG 2014, PRD 2016)
- NR Surrogates: Blackman et al
- BNS models: Lackey et al (PRD 2017)
- Reduced-order quadratures: Smith+ (PRD 2016), Canizares+ (PRL 2015), Antil+ (JSC 2013)

Reduced order (surrogate) model schematic

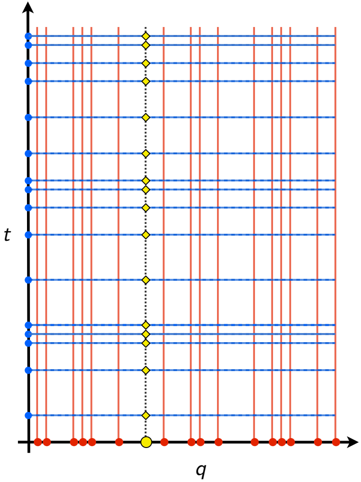

1. Create training dataset from N model evaluations

Reduced order (surrogate) model schematic

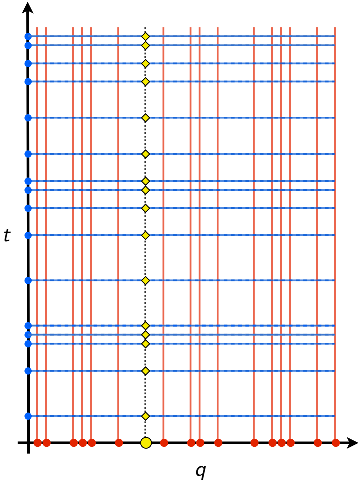

1. Create training dataset from N model evaluations
2. Compress the model with $n \leq N$ model-specific basis set $h_i^{\mathrm{basis}}(t)$

Reduced order (surrogate) model schematic

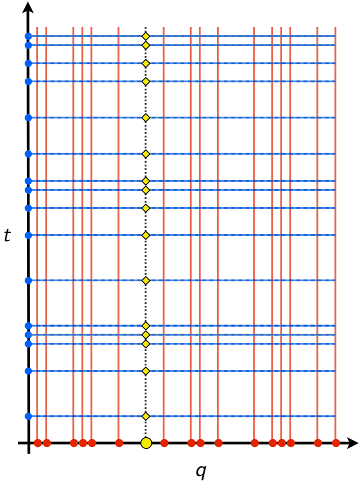

1. Create training dataset from N model evaluations
2. Compress the model with $n \leq N$ model-specific basis set $h_i^{\mathrm{basis}}(t)$
3. n specially selected times T_i

Reduced order (surrogate) model schematic

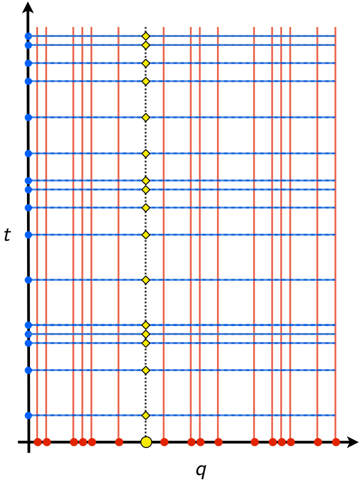

1. Create training dataset from N model evaluations
2. Compress the model with $n \leq N$ model-specific basis set $h_i^{\mathrm{basis}}(t)$
3. n specially selected times T_i
4. Parametric fits $h_\mu^{\mathrm{FIT}}(q; T_i)$ at each T_i

Reduced order (surrogate) model schematic

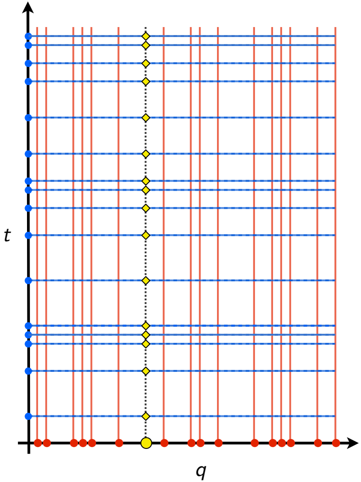

1. Create training dataset from N model evaluations
2. Compress the model with $n \leq N$ model-specific basis set $h_i^{\mathrm{basis}}(t)$
3. n specially selected times T_i
4. Parametric fits $h_\mu^{\mathrm{FIT}}(q; T_i)$ at each T_i
5. Evaluate the surrogate at parameter value (yellow dot) by i) evaluating the fits at each T_i which ii) specifies the full waveform through an (empirical) interpolant:

$$h_\mu^{\mathrm{S}}(t; q) \equiv \sum_{i=1}^{m} B_i(t) h_\mu^{\mathrm{FIT}}(q; T_i)$$

where $\{B_i\}$ is built from $h_i^{\mathrm{basis}}(t)$

Reduced order (surrogate) model schematic – another view

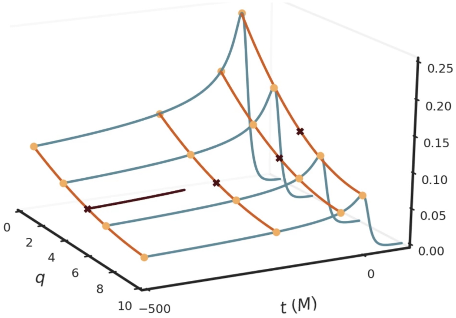

Fig: Vijay Varma

1. Create training dataset from N model evaluations
2. Compress the model with $n \leq N$ model-specific basis set 3. n specially selected times T_i
4. Parametric fits $h_\mu^{\mathrm{FIT}}(q; T_i)$ at each T_i
5. Evaluate the surrogate at parameter value (yellow dot) by i) evaluating the fits at each T_i which ii) specifies the full waveform through an (empirical) interpolant:

$$h_\mu^{\mathrm{S}}(t; q) \equiv \sum_{i=1}^{m} B_i(t) h_\mu^{\mathrm{FIT}}(q; T_i)$$

where $\{B_i\}$ is built from $h_i^{\mathrm{basis}}(t)$

Step 1: basis generation

- Seek a representation of the gravitational wave model

$$h_\mu(t) \approx \sum_{i=1}^{m} c_i(\mu) e_i(t) \qquad \text{or} \qquad h_\mu(f) \approx \sum_{i=1}^{m} c_i(\mu) e_i(f)$$

for m as small as possible and $\mu = (\mathrm{mass}, \mathrm{spin}, \dots)$ labels the parameterization

- Sometimes referred to as a *reduced order model* (model is *reduced* to m degrees of freedom)

Whats special about e_i ???

- Application-specific basis
- Fewer basis $\rightarrow$ faster computations

Some methods: Greedy-RB and singular value decomposition algorithms (details later).

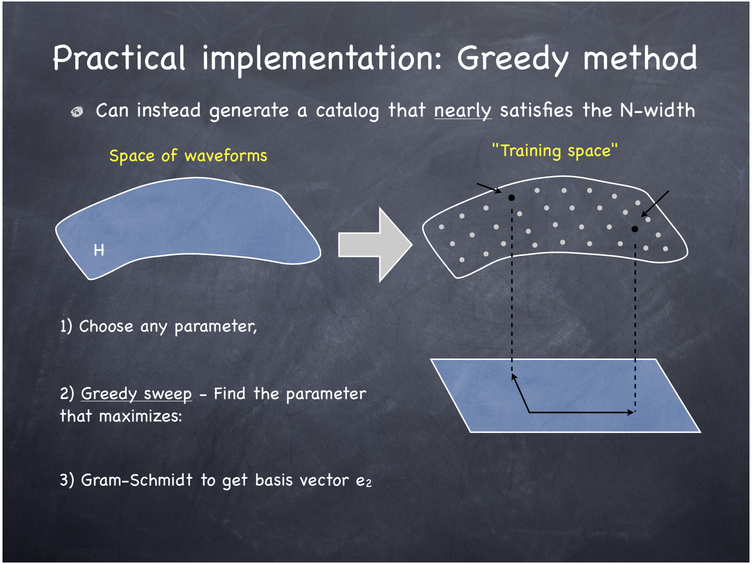

Slide courtesy of Chad Galley

Example

Effective one body (Pan et al., 2011)

- (2,2) mode for $q \in [1,2]$, duration $\approx 12,000$M
- Fast decay of approximation (overlap) error

$$\max_q \|h_q - \sum_{i=1}^m c_i(q) e_i\|^2$$

Other evidence

- Observed across models, regimes
- Observed by groups using POD/SVD
    - Cannon et al (PRD 044025)
    - M. Püerrer (arXiv:1402.4146)

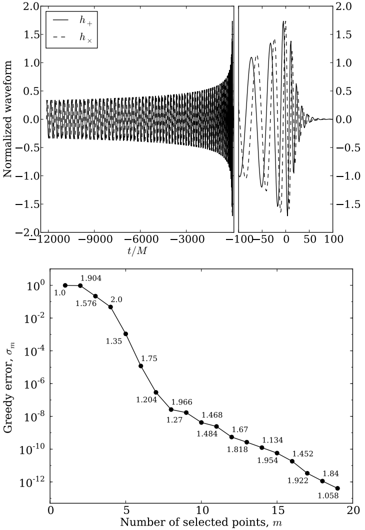

Waveform compression application (ex: $q \sim 1.2040$)

Ortho. Basis

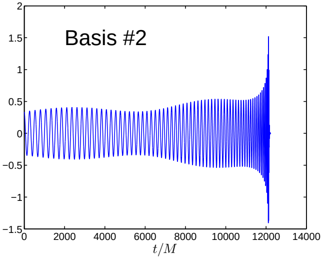

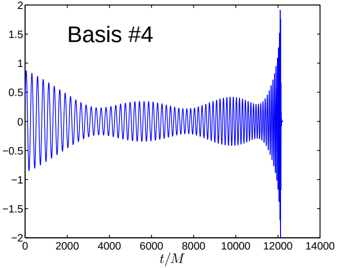

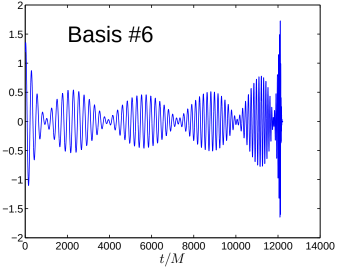

Approx:

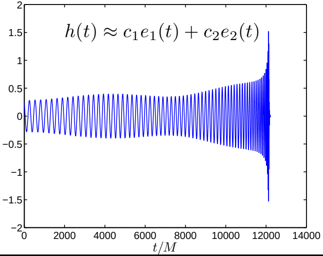

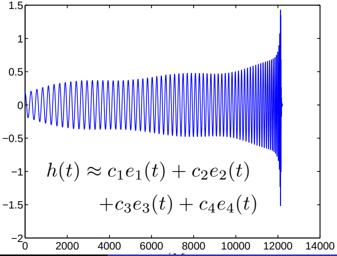

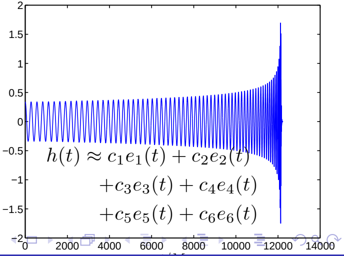

Example: Parameterized Heaviside (toy IMR model)

Continuum:

$$H(\mu - x)$$

$$x \in [-1, 1]$$

$$\mu \in [-.2, .2]$$

Training set:

$$\{H(\mu_i - x)\}$$

$$\mu_i = -.2 + \frac{.4}{4000} i$$

$$i \in [0, \dots, 4000]$$

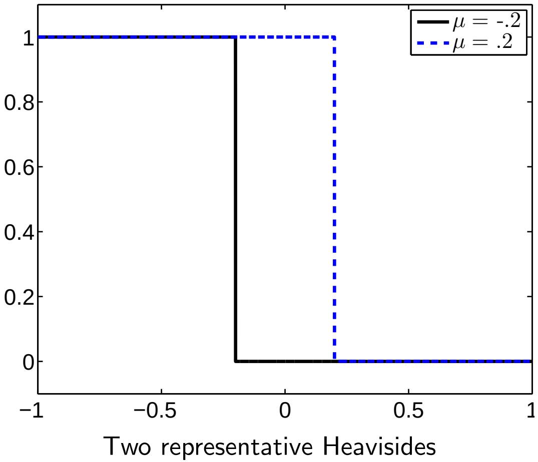

Example: Parameterized Heaviside (toy IMR model)

1. Select first basis (seed):

$$H(-.2 - x)$$

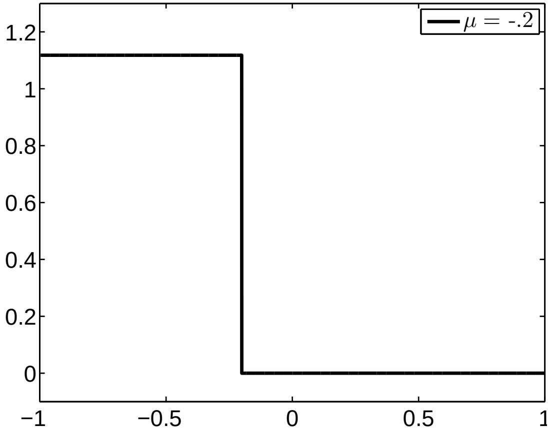

Example: Parameterized Heaviside (toy IMR model)

1. Select first basis (seed):
$$H(-.2-x)$$

2. Find worst approximation:
$$\textcolor{blue}{\mathrm{Err}_i} =$$
$$H(\mu_i - x) - \textcolor{red}{c}H(-.2-x)$$

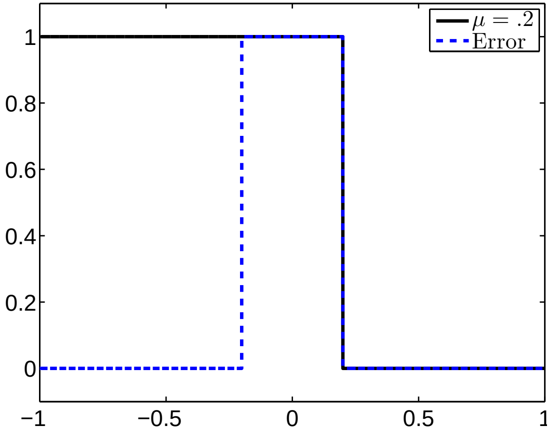

Example: Parameterized Heaviside (toy IMR model)

1. Select first basis (seed):
$$H(-.2 - x)$$
2. Find worst approximation:
$$\mathrm{Err}_i = H(\mu_i - x) - cH(-.2 - x)$$
3. Second basis:
$$\mu = .2 \quad \rightarrow \quad H(.2 - x)$$

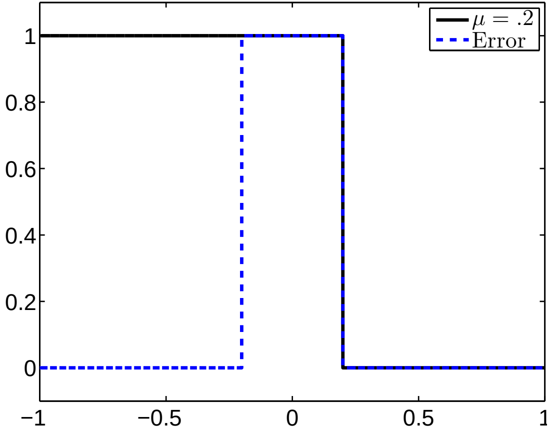

Repeat steps 2 & 3 until an approximation threshold is achieved

Greedy output (basis):

$$\mu^{\mathrm{RB}} = \{ -0.2, 0.2, 0.0203, -0.0844, 0.1147, \ldots \}$$

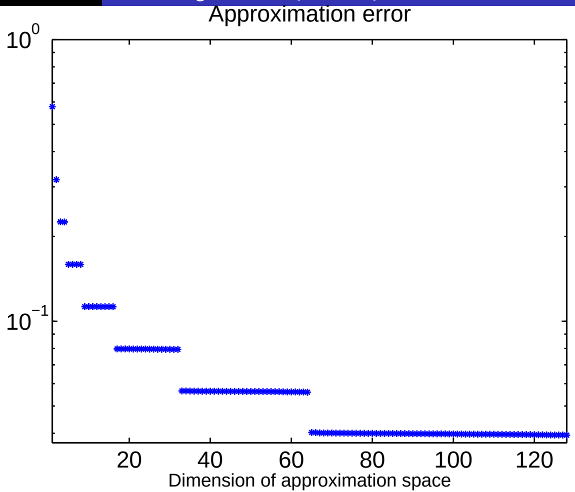

- Greedy algorithm "fails" (SVD will too). Non-smooth w.r.t. parameter variations.
- If we let $y(\mu) = \mu - x$ then only 1 basis function $H(y)$ needed

- Need a fast/accurate way to compute the coefficients for *any* parameter μ

$$h_\mu(t) \approx \sum_{i=1}^{m} c_i(\mu) e_i(t)$$

- A convenient expression for $c_i(\mu)$ thanks to approximation theory...

- Need a fast/accurate way to compute the coefficients for *any* parameter μ

$$h_\mu(t) \approx \sum_{i=1}^m c_i(\mu) e_i(t)$$

- A convenient expression for $c_i(\mu)$ thanks to approximation theory...

Given m basis, there (usually) exists m times $\{T_i\}_{i=1}^m$ for which

$$\{c_i(\mu)\}_{i=1}^m \Longleftrightarrow \{h_\mu(T_i)\}_{i=1}^m$$

m numbers c_i contain equivalent information as m numbers $h_\mu(T_i)$

Empirical interpolation method

- Input: m basis $\{e_i(t)\}_{i=1}^{m}$
- Output: Nearly optimal selection of m times $\{T_i\}_{i=1}^{m}$
- These times are adapted to the problem/basis - unlike Chebyshev nodes

Barrault 2004, Maday 2009, Chaturantabut 2009, Sorensen 2009

Interpolation points for EOB waveforms

What are the best temporal interpolation points for an EOB-basis?

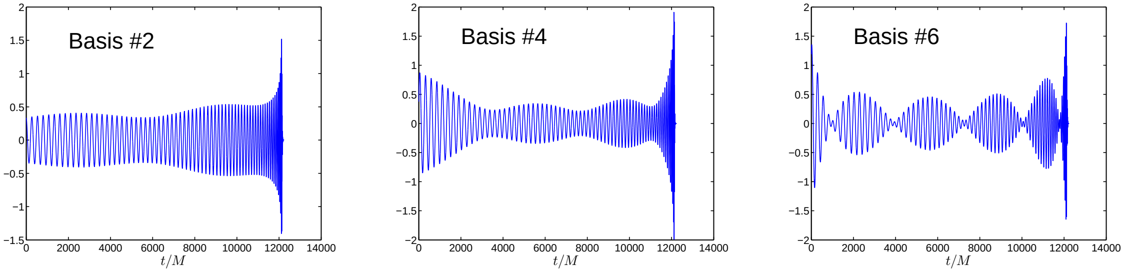

Model: non-spinning EOB, $q \in [1, 2]$, 65-70 wave cycles (previous example)

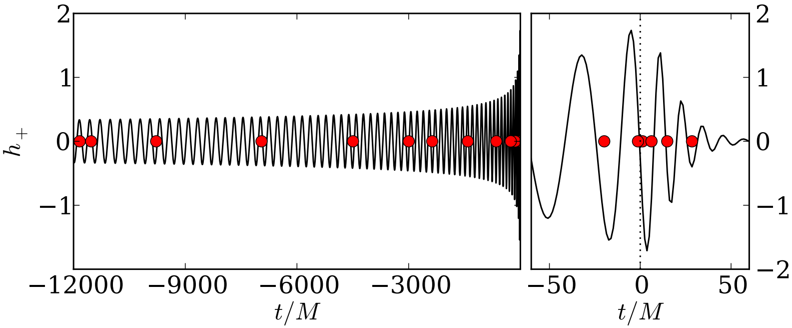

- Any waveform in the above range can be recovered through its evaluation at these 5 (error $\sim 10^{-4}$) to 19 (error $\sim 10^{-12}$) empirical time nodes

$$h_\mu(t) \approx \sum_i^m c_i(\mu) e_i(t) = \sum_i^m B_i(t) h_\mu(T_i)$$

Step 3: parametric fits

Main idea: Make the data look boring

- We know how $h_\mu(T_i)$ should look

$$h_\mu(t) \equiv A_\mu(t)\mathrm{e}^{-\mathrm{i}\phi_\mu(t)}$$

- A and ϕ are boring!
- Polynomial fit (in q) works well

$$A_q(T_i) \approx A_i(q) = \sum_{n=0}^{N} a_n q^n$$

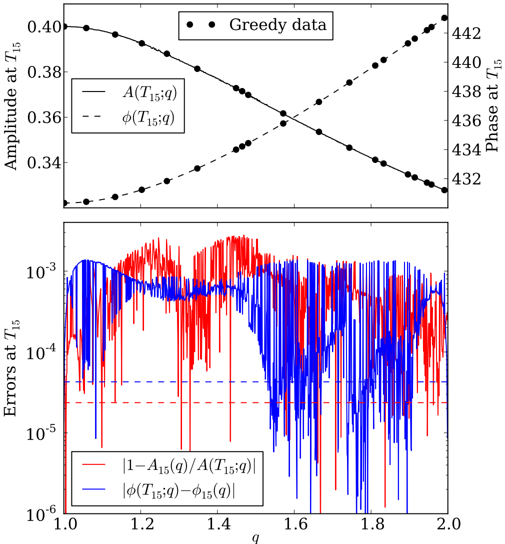

Alternative choices

$$h_\mu(t) \approx \sum_i^m c_i(\mu) e_i(t)$$

... no amplitude and phase decomposition.

$c_i(\mu) \rightarrow$

- $c_i(\mu)$ "exotic" looking function
- Deciding form of data to approximate is important ("feature engineering")

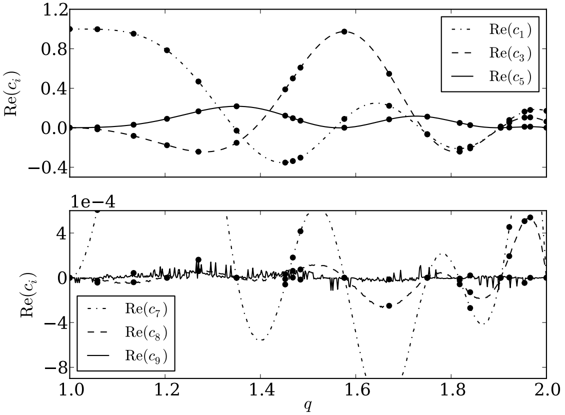

c_i vs q

Summary

These three (offline) steps complete the building phase

Final step: validate the model

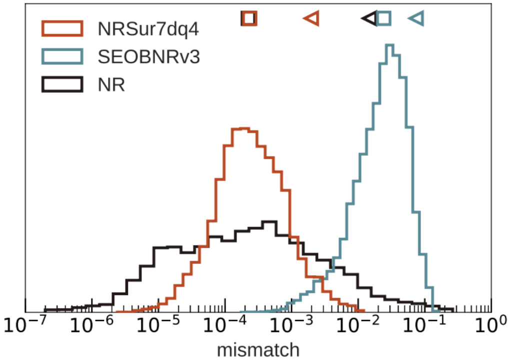

Cross-validation (arXiv:1905.09300).

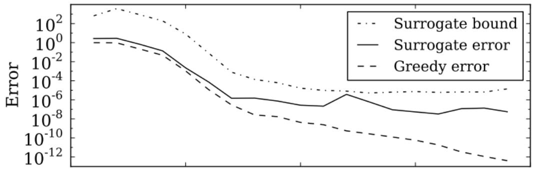

Error bounds (arXiv:1308.3565)

Computing lab

Jupyter notebook demo (Today): building a 1-dimensional (2,2)-mode IMR model

Going further (Homework): Higher-dimensional models require more complicated data decomposition and regression tools suited for high-dimensional data on scattered grids.