

TLDR: time lag/delay reconstructor

M. D. Anderson¹, F. Baron^{1,2} and M. C. Bentz¹

¹Georgia State University, 25 Park Place NE, Atlanta, GA 30303-2911, USA

²Center for High Angular Resolution Astronomy, P.O. Box 5060, Atlanta, GA 30302-5060, USA

Accepted 2021 May 6. Received 2021 May 6; in original form 2020 November 2

ABSTRACT

We present the time lag/delay reconstructor (TLDR), an algorithm for reconstructing velocity delay maps in the maximum a posteriori framework for reverberation mapping. Reverberation mapping is a tomographical method for studying the kinematics and geometry of the broad-line region of active galactic nuclei at high spatial resolution. Leveraging modern image reconstruction techniques, including total variation and compressed sensing, TLDR applies multiple regularization schemes to reconstruct velocity delay maps using the alternating direction method of multipliers. Along with the detailed description of the TLDR algorithm we present test reconstructions from TLDR applied to synthetic reverberation mapping spectra as well as a preliminary reconstruction of the $H\beta$ feature of Arp 151 from the 2008 Lick Active Galactic Nuclei Monitoring Project.

Key words: techniques: image processing – galaxies: active – galaxies: nuclei.

1 INTRODUCTION

In an active galactic nucleus (AGN), clouds of photoionized gas near the supermassive black hole (SMBH) at the galaxy’s centre emit highly broadened emission lines due to the extreme kinematics of the region. The area in which these broad lines are emitted is known as the broad-line region (BLR). Rapid variation of the broad-line emission in response to variations in the continuum emission has long been observed for both radio quiet and radio loud AGNs (Sandage 1967; Cromwell & Weymann 1970; Penston et al. 1974; Peterson 1988; Ulrich, Maraschi & Urry 1997).

While the precise mechanism behind the observed variability is unknown, a number of viable theories exist, including, but not limited to, simple accretion on to the SMBH, magnetic field interactions within the accretion disc, or some combination of processes varying in contribution over various time-scales (Czerny 2006). Regardless of the mechanism at play, the release of energy is observed as continuum light that propagates outwards from the central region of the AGN. During the process, clouds of gas are ionized at different times depending on their distance from the accretion disc. As the photoionized gas recombines, it re-emits the light as an emission line feature subject to the kinematics of the cloud. This is to say that any observed change above the noise in continuum emission from the AGN will be followed by a change in emission lines in the BLR separated by a delay corresponding to the difference in path-length between the directly observed continuum emission, the resulting emission line feature, and the observer.

The classic example is that of a thin ring as described in Peterson (1993), shown in schematic form in Fig. 1, where the observer lies in plane of the figure at a sufficient distance along the horizontal line that light traveling from the thin ring to the observer arrive in effectively flat wavefronts. The thin ring of gas at a radius of one

light-day from the central black hole is ionized by a pulse of light described by a δ -function. The ionizing pulse reaches the gas that comprises the shell at the same time. The gas is ionized and the energy is then re-emitted by the gas, subject to Doppler effects due to the kinematics of the region. Some of this re-emitted light is along the line of sight to the observer. Light that is re-emitted at the far side of the ring must travel back across the region, a distance of twice the ring’s radius, such that it arrives at the observer with a delay of $2r/c$. Following the arc of the ring from the near side to the far side, the observer will see the re-emitted light arrive with delay times covering the range from 0 to $2r/c$ with a bulk average delay of r/c . In Fig. 1, the dashed lines show lines of constant delay, known as isodelay surfaces, across the thin ring at varying delay times τ .

Moving from the thin ring to a more realistic distribution of gas around the SMBH results in more complex patterns in the observed response. Reverberation mapping (RM) is the tomographical method used to study the arrangement of photoionized gas near accreting black holes (Blandford & McKee 1982). At present, RM is primarily used to measure the masses of the SMBH at the heart of AGNs (Peterson & Wandel 2000; Bentz et al. 2009; Bentz & Katz 2015; De Rosa et al. 2015; Shen et al. 2015). RM results also establish the radius–luminosity relationship (Bentz 2009; Bentz et al. 2013), which allows us to estimate SMBH masses across all redshifts and consequently consider their influence on the evolution of galaxies across cosmological time (Booth & Schaye 2009).

Studying the kinematics of AGNs at the geometric scales provided by RM represents a substantial impetus behind RM campaigns (Peterson 1993; Horne et al. 2004). By providing complementary information on scales smaller than may be probed via integral field spectroscopy (Storchi-Bergmann & Schnorr-Müller 2019) or interferometry (Gravity Collaboration 2020), RM has the potential to bridge the gap in spatial scale between inflows and outflows and the central engine of AGNs.

One goal of RM in the inverse problem approach is to recover the velocity delay map (VDM) of an observed AGN. The VDM is

* E-mail: manderson@astro.gsu.edu




Figure 1. Classic isodelay surface schematic of an AGN showing how the various delays sample the entire area of the AGN consisting of a thin ring of spectral emitting material about the central SMBH. Based on Figure 1 from Peterson (1993), with permission.

a representation of the BLR where the geometry is encoded in the observed delay time and line-of-sight velocity of the re-emitting gas. Mathematical recovery of the VDM results from solving an ill-posed inverse problem, meaning that the number of degrees of freedom to reconstruct (the VDM) is greater than the effective number of data points (flux-varying spectra). A conventional solution to this problem is regularized maximum likelihood or maximum a posteriori (MAP) in a Bayesian framework. Most software packages for RM rely on this formalism (Horne 1994; Krolik 1994; Skielboe et al. 2015), where the χ^2 of the reconstructed map is minimized along a set of regularizers that enforce the prior expectations on the model, the exception being a direct modelling approach (Pancoast, Brewer & Treu 2014).

Classic examples of regularizers are positivity and maximum entropy (to enforce smoothness). Note that without these priors, the reconstruction algorithm would overfit the χ^2 , giving rise to spurious artefacts which would make the VDM uninterpretable. On the other hand, regularizers must remain generic, i.e. as non-committal as possible, so as not to bias the reconstruction towards a non-physical solution. As shown in other fields of study such as interferometry (Thiebaud & Giovannelli 2010; Schutz et al. 2015) and medical imaging (Cao, Nehorai & Jacob 2007; Yao et al. 2015; Li et al. 2016), it is the quality of the regularization and its strength that determine the actual mapping fidelity of a given piece of software. Unfortunately, currently available software packages for RM offer little to no flexibility in terms of regularizers, and are mostly centred on smoothing, which *de facto* limit the interpretation of the data. From recent advances in the theory of Compressed Sensing, we now know that the key to optimally regularize such inverse problems is tapping into the sparsity of the reconstructed signals (Candes & Wakin 2008). In image reconstruction, sparsity is the idea that a signal may be represented by a number of non-zero coefficients that is small relative to the total number of coefficients in the signal (Davenport et al. 2011; Bryan & Leise 2013). RM is a poster example for Compressed Sensing, in that the VDMs that have been previously recovered for various AGNs are sparse.

The time lag/delay reconstructor (TLDR) algorithm presented here utilizes the alternating direction method of multipliers (ADMM),

as well as two-dimensional regularization across the temporal and spectral axes of RM data, to reconstruct VDMs for RM.

This paper serves as an introduction to image reconstruction in the MAP framework as applied to RM. It will discuss the basics of RM in Section 2, the inverse problem approach and appropriate regularizers in Section 3, the basics of ADMM in Section 4, the construction of the TLDR algorithm as an application of ADMM in Section 5, basic reconstruction tests in Section 6, and offer a preliminary reconstruction of the Arp 151 $H\beta$ feature in Section 7.

2 REVERBERATION MAPPING

RM is a tomographical imaging process in which rapid changes in continuum emission from the centre of an AGN is used to map the kinematics and geometry of the gas that re-emits the continuum emission as line emission. While the precise relationship between the ionizing and emitted radiation is unknown, a number of assumptions serve as the backbone of RM. These assumptions as laid down in Peterson (1993) are:

- (i) The continuum emission originates from a compact isotropically emitting central source.
- (ii) The light-traveltime between the central source and the clouds of the BLR is the primary source of delay.
- (iii) A simple relationship between the ionizing emission of the central source and the response observed in BLR exists. This relationship need not be strictly linear in nature.

In reality, continuum emission does not appear in δ -function pulses, but rather as continuous variations. Both the spectral line and continuum emission must be treated as continuous functions. Doing so, the observed line emission, $\mathbf{L}(t, \lambda)$ and continuum emission, $\mathbf{C}(t)$, are functions of time which are related by the VDM \mathbf{x} (Peterson 1994). Mathematically, the emission line response is the result of the continuum being convolved with the VDM, also called the transfer function (Blandford & McKee 1982),

$$\mathbf{L}(t, \lambda) = \int_0^{\infty} \mathbf{x}(\tau, \lambda) \mathbf{C}(t - \tau) d\tau. \quad (1)$$

Due to the discrete nature of observations, the convolution is not continuous, but discrete which appears as

$$\mathbf{L}[t, \lambda] = \sum_{\tau=0}^N \mathbf{x}[\tau, \lambda] \mathbf{C}[t - \tau]. \quad (2)$$

The nature of the element-wise convolution allows the construction of a mapping matrix \mathbf{H} comprised of elements from the continuum so that equation (2) can be represented purely as the linear relation

$$\mathbf{L} = \mathbf{H}\mathbf{x}. \quad (3)$$

The matrix \mathbf{H} is constructed with rows corresponding to the number of observations in the spectra taken at time t and columns corresponding to the delay times τ such that each element in the matrix is a value of the continuum flux measured at time $t - \tau$. The resulting mapping matrix \mathbf{H} is a Toeplitz matrix representing the one-dimensional convolution. Since the entire problem is linear, recovering the VDM is a straightforward inverse problem.

It is evident in equation (3) that this method requires that the continuum data must exist at the proper sample times $t - \tau$ for all spectral sample dates t and all delay times τ . The required high cadence represents a serious challenge for all RM observational campaigns. In earlier work on RM, the requirement that continuum observations exist on a grid has been met by linearly interpolating the continuum

data (Horne 1994; Krolik 1994). Recently, it has become common to model the continuum data in some way where both damped random walk (Kelly, Bechtold & Siemiginowska 2009; Zu, Kochanek & Peterson 2011) and Gaussian process (Pancoast, Brewer & Treu 2011) models have been used to interpolate continuum observations.

3 INVERSE PROBLEM APPROACH

In order to find the VDM, we adopt here the general framework for MAP. Prior knowledge about the system, such as positivity, smoothness, etc., is introduced in the form of regularization terms (Thiébaud 2005) that are minimized alongside the neg-log-likelihood of the data. The VDM solution \mathbf{x}_{opt} is given by

$$\mathbf{x}_{\text{opt}} = \arg \min_{\mathbf{x} \in \mathbb{R}^{M \times N}} \{f_{\text{data}}(\mathbf{x}) + f_{\text{prior}}(\mathbf{x})\}. \quad (4)$$

Assuming Gaussian statistics based on flux levels, exposure times, and the Central Limit Theorem, the f_{data} term is commonly known as the χ^2 of the model relative to the data. In the case that Gaussian statistics do not hold and Poisson noise dominates, ADMM could still be used. Here, construction of the χ^2 involves combining the discrete convolution of equation (3) with knowledge of the error in the data,

$$f_{\text{data}}(\mathbf{x}) = \frac{1}{2} \chi^2(\mathbf{x}) = \frac{1}{2} \|\mathbf{L} - \mathbf{H}\mathbf{x}\|_{\mathbf{W}}^2 = \frac{1}{2} (\mathbf{H}\mathbf{x} - \mathbf{L})^{\top} \mathbf{W} (\mathbf{H}\mathbf{x} - \mathbf{L}), \quad (5)$$

where \mathbf{W} is the inverse covariance matrix of the emission line data \mathbf{L} .

The second term in equation (4) is the regularization term where the prior knowledge and expectations of the problem are enforced. The regularization term consists of the weighted sum of all regularizers, so that f_{prior} can be written as

$$f_{\text{prior}}(\mathbf{x}) = \sum_i \mu_i R_i(\mathbf{x}). \quad (6)$$

The combination of each regularizer, R_i , and its associated regularization hyperparameter, μ_i , allows each regularization term to be weighted individually allowing for flexibility in the reconstruction process. These functions serve to penalize the likelihood term in the minimization equation (equation 4). The effect of these penalties is to increase the value of the function within the minimizer, making these solutions less probable in the Bayesian framework. Selection of regularizers for a reconstruction is an important step in setting up the reconstruction. Aside from being important to the eventual solution the use of different regularization strategies may impact the problem to be fed into the minimization algorithm.

For the RM problem, there are a number of regularizations that are natural choices. Foremost among these is positivity. Due to the nature of the AGN system creating the continuum pulses and the linear relationship between the pulses and the re-emission from the reverberating medium, only positive delays have any physical meaning. Negative values can only arise in the reconstruction where the linear model does not hold. The positivity regularizer works as a filter where negative values are simply projected to zero.

Since previously reconstructed VDMs are visibly sparse in the spectral axis (Bentz et al. 2010; Grier et al. 2013), another obvious regularizer choice is that of the ℓ_1 -norm. More specifically, in the spectrum of an AGN at wavelengths where no broad emission lines are present, the VDM is expected to be zero for all delay times. As previously mentioned, this sparsity makes the RM problem a particularly good example for compressed sensing. As such, the so-called ℓ_1 -norm is well suited for reconstructing VDMs. The ℓ_1 -norm

favours the fewest number of non-zero terms in \mathbf{x} . In the case of RM this means that solutions with more zero-valued VDMs will be favoured over any other solutions, matching with the expected sparsity in the VDM. Regularization by the ℓ_1 -norm is accomplished according to the soft thresholding scheme of Donoho (1995) with an additional enforcement of positivity.

As an alternative to the ℓ_1 -norm, the square of the ℓ_2 -norm favours images with many zero valued coefficients. It is however, slightly different mathematically and its primary purpose is to promote smoothness in the image. The ℓ_2 -norm penalizes larger pixel values more than small ones due to the square of the values seen in equation. Reducing the value of the ℓ_2 -norm-squared drives the total amount of flux in the image down and brings the maximum and minimum flux values closer together, promoting smoother images.

Perhaps most importantly, the VDM is expected to be smooth both along the temporal axis and the spectral axis. The temporal axis is expected to be smooth due to the relationship between absorbed and emitted light from the gas clouds being mapped in the BLR. The spectral axis will be smooth due to the kinematics of the region and the Doppler broadening of the emission lines themselves. Compressed sensing theory prescribes to enforce smoothness based on sparsity, either that of the gradient of the solution, or even more generally that of its wavelet transform. In both cases, regularization along the different axes is better treated as two individual regularization steps, allowing the regularization along each axis to be weighted differently.

In the case of gradient sparsity, the regularizer is known as anisotropic total variation (TV) and given by the ℓ_1 -norm of the gradient of the image (Rudin, Osher & Fatemi 1992). Minimizing this regularization term favours \mathbf{x} such that the sum over the gradient is as small as possible, which is to say the solutions with the least variations are favoured. TV in this case is constructed using a finite difference matrix operator to represent the gradient, and the TV regularization term can be easily calculated and minimized using the same Soft Thresholding scheme as the standard ℓ_1 -norm, albeit without the enforcement of positivity.

4 ADMM

In order to solve the problem posed by equation (12), ADMM is used. ADMM was developed as a combination of the method of multipliers and the dual ascent method (Boyd et al. 2011; Eckstein & Yao 2012). The convergence behaviour of ADMM is well documented in general (Nishihara et al. 2015; Wang, Yin & Zeng 2018) and for linear and quadratic problems specifically (Boley 2013; Han & Yuan 2013). This section is comprised of a discussion of ADMM for global variable consensus optimization (GVCO) and how it solves problems of the form of equation (12).

GVCO ADMM seeks to find the solution to the general GVCO problem,

$$\arg \min_{\mathbf{x}_i} \sum_{i=1}^N f_i(\mathbf{x}_i) \quad (7)$$

subject to $\mathbf{x}_i - \mathbf{z} = 0, i = 1, \dots, N.$

While having a different form than the general equation for solving an inverse problem in the MAP framework (equation 4), equation (7) represents the same problem. The differences lie in the construction of the ADMM algorithm. Key in the construction of ADMM is a variable splitting. In ADMM, the problem must be separable so that the general function can be split into two groups of functions, $f_i(\mathbf{x}_i)$ and $g(\mathbf{z})$. ADMM will then seek to bring these two function into consensus so that at convergence $\mathbf{x}_i = \mathbf{z}$.

Table 1. ADMM variable assignments within TLDR.

Regularizer	Original term	Variable assignment	New term	Penalty parameter	Lagrange multiplier
Smoothing	$\ \mathbf{H}\mathbf{x} - \mathbf{L}\ _2^2 + \frac{\mu_{\ell_2}}{2} \ \mathbf{x}\ _2^2$	$\mathbf{x}_0 = \mathbf{X}$	$\ \mathbf{H}\mathbf{X} - \mathbf{L}\ _F^2 + \frac{\mu_{\ell_2}}{2} \ \mathbf{X}\ _2^2$	ρ_X	\mathbf{Y}_X
Positivity	$\mathbb{1}_{\mathbb{R}^+}(\mathbf{x})$	$\mathbf{x}_1 = \mathbf{P}$	$\mathbb{1}_{\mathbb{R}^+}(\mathbf{P})$	ρ_P	\mathbf{Y}_P
ℓ_1 -norm	$\mu_{\ell_1} \ \mathbf{x}\ _1$	$\mathbf{x}_2 = \mathbf{N}$	$\mu_{\ell_1} \ \mathbf{N}\ _1$	ρ_N	\mathbf{Y}_N
Total variation	$\mu_T \ \nabla \mathbf{x}\ _1$	$\mathbf{x}_3 = \nabla \mathbf{Z} = \mathbf{T}$	$\mu_T \ \mathbf{T}\ _1$	ρ_T	\mathbf{Y}_T

The ADMM algorithm is an augmented Lagrangian method, so the algorithm relies on the augmented Lagrangian in its minimization steps. The augmented Lagrangian takes the general form,

$$\mathcal{L}(\mathbf{x}_1, \dots, \mathbf{x}_N, \mathbf{z}, \mathbf{y}) = \sum_{i=1}^N \left(f_i(\mathbf{x}_i) + \mathbf{y}_i^\top (\mathbf{x}_i - \mathbf{z}) + \frac{\rho_i}{2} \|\mathbf{x}_i - \mathbf{z}\|_2^2 \right) \quad (8)$$

where \mathbf{y} represents the Lagrange multipliers and ρ the penalty parameter.

ADMM then proceeds iterating and following these three steps, where the current iteration number is denoted by a superscript k :

(i) Minimize each term, subscript i , of the augmented Lagrangian with respect to the corresponding \mathbf{x}_i ,

$$\mathbf{x}_i^{k+1} = \arg \min_{\mathbf{x}_i} \left(f_i(\mathbf{x}_i^k) + \mathbf{y}_i^{\top, k} (\mathbf{x}_i^k - \mathbf{z}^k) + \frac{\rho_i}{2} \|\mathbf{x}_i^k - \mathbf{z}^k\|_2^2 \right). \quad (9)$$

(ii) Minimize the augmented Lagrangian with respect to \mathbf{z} ,

$$\mathbf{z}^{k+1} = \arg \min_{\mathbf{z}} \left(\sum_{i=1}^N -\mathbf{y}_i^{\top, k} \mathbf{z}^k + \frac{\rho_i}{2} \|\mathbf{x}_i^{k+1} - \mathbf{z}^k\|_2^2 \right). \quad (10)$$

(iii) Update the Lagrange multipliers,

$$\mathbf{y}_i^{k+1} = \mathbf{y}_i^k + \rho_i (\mathbf{x}_i^{k+1} - \mathbf{z}^{k+1}). \quad (11)$$

This procedure repeats until whatever convergence parameters are in place are satisfied. This method offers flexibility in that numerous regularization terms can be added with relatively little difficulty. The method also offers the potential for great computational speed in that many steps and subroutines in the algorithm can be implemented in parallel.

5 TLDR

Inspired by recent implementations of ADMM in interferometry (Thiébaud & Soulez 2012; Schutz et al. 2014; Ferrari et al. 2015), the TLDR algorithm is built as an implementation of GVCO ADMM constructed to solve the RM problem given by equation (12). Through the assignment of new variable names, one for each regularization term used, equations (9)–(11) can be expanded to build the algorithm. Within TLDR, the χ^2 and ℓ_2 -norm smoothing regularizations operate on \mathbf{X} , the positivity regularization operates on \mathbf{P} , the compressed sensing regularizer operates on \mathbf{N} , and the total variation regularizer operates on \mathbf{T} . The new variables and their associated penalty parameters and Lagrange multipliers help to distinguish mathematical operations in the code and are listed in Table 1. For convenience, the \mathbf{z} variable is assigned to \mathbf{Z} .

With the new variables in place, the problem can be rewritten as follows:

$$\begin{aligned} & \arg \min_{\mathbf{X}} \frac{1}{2} (\mathbf{H}\mathbf{X} - \mathbf{L})^\top \mathbf{W} (\mathbf{H}\mathbf{X} - \mathbf{L}) + \frac{\mu_X}{2} \|\mathbf{X}\|_2^2 + \\ & \mu_N \|\mathbf{N}\|_1 + \mathbb{1}_{\mathbb{R}^+}(\mathbf{P}) + \mu_T \|\mathbf{T}\|_1 \\ & \text{subject to } \mathbf{X}, \mathbf{P}, \mathbf{N} = \mathbf{Z}; \mathbf{T} = \nabla \mathbf{Z}. \end{aligned} \quad (12)$$

Expanding equation (8) using the new variables and the problem as written in equation (12) allows us to easily write the full augmented Lagrangian:

$$\begin{aligned} \mathcal{L} = & (\mathbf{H}\mathbf{X} - \mathbf{L})^\top \mathbf{W} (\mathbf{H}\mathbf{X} - \mathbf{L}) + \frac{\mu_{\ell_2}}{2} \|\mathbf{X}\|_2^2 \\ & + \mathbf{Y}_X^\top (\mathbf{X} - \mathbf{Z}) + \frac{\rho_X}{2} \|\mathbf{X} - \mathbf{Z}\|_2^2 \\ & + \mathbb{1}_{\mathbb{R}^+}(\mathbf{P}) + \mathbf{Y}_P^\top (\mathbf{P} - \mathbf{Z}) + \frac{\rho_P}{2} \|\mathbf{P} - \mathbf{Z}\|_2^2 \\ & + \mu_{\ell_1} \|\mathbf{N}\|_1 + \mathbf{Y}_N^\top (\mathbf{N} - \mathbf{Z}) + \frac{\rho_N}{2} \|\mathbf{N} - \mathbf{Z}\|_2^2 \\ & + \mu_T \|\mathbf{T}\|_1 + \mathbf{Y}_T^\top (\mathbf{T} - \nabla \mathbf{Z}) + \frac{\rho_T}{2} \|\mathbf{T} - \nabla \mathbf{Z}\|_2^2. \end{aligned} \quad (13)$$

TLDR proceeds through the steps of ADMM outlined in Section 4 shown as pseudo-code in Algorithm 1.

Algorithm 1: Pseudocode for TLDR's ADMM block.

```

initialization  $\mathbf{X}, \mathbf{P}, \mathbf{N} = \mathbf{Z}; \mathbf{T} = \nabla \mathbf{Z}$ 
calculate Initial VDM
while converged == False do
    minimization of Augmented Lagrangian w.r.t.  $\mathbf{X}$ 
    minimization of Augmented Lagrangian w.r.t.
        regularization terms  $\mathbf{P}, \mathbf{N}, \mathbf{T}$ 
    minimization of Augmented Lagrangian w.r.t.  $\mathbf{Z}$ 
    update Lagrange Multipliers  $\mathbf{Y}_X, \mathbf{Y}_P, \mathbf{Y}_N, \mathbf{Y}_T$ 
    check for convergence
end
    
```

Aside from the Lagrange multipliers update, which is performed exactly as equation (11), the steps in Algorithm 1 are non-trivial and each warrants some discussion. So each will be discussed in order.

5.1 Initial VDM

The TLDR algorithm could begin from any initial VDM, but the optimal starting delay map should be as close to the final solution as possible. This reduces the amount of parameter space over which the pixel values in the delay map must move to find the best values. By starting near to the final solution, the number of iterations required by the algorithm is reduced thereby reducing the amount of time required for the algorithm to find the final solution. However, care must be taken as it is possible to introduce artefacts into the initial VDM that will not easily be removed within the rest of the algorithm.

By isolating the first term in the augmented Lagrangian, a singly regularized least-squares equation for the RM problem is found:

$$f(\mathbf{X}) = (\mathbf{H}\mathbf{X} - \mathbf{L})^\top \mathbf{W} (\mathbf{H}\mathbf{X} - \mathbf{L}) + \frac{\mu_{\ell_2}}{2} \|\mathbf{X}\|_2^2. \quad (14)$$

Minimizing, by taking the derivative of the function $f(\mathbf{X})$, setting it equal to zero, and solving for \mathbf{X} yields the initial VDM:

$$\mathbf{X}^0 = (\mathbf{H}^\top \mathbf{W} \mathbf{H} + \mu_{\ell_2} \mathbf{I})^{-1} (\mathbf{H}^\top \mathbf{W} \mathbf{L}). \quad (15)$$

This method for analytically solving the ℓ_2 -squared regularized least-squares problem is known as the Tikhonov regularization (Tikhonov 1963) or ridge regression (Hoerl & Kennard 1970).

5.2 Minimization with respect to \mathbf{X}

The minimization for \mathbf{X} can be found analytically by solving the derivative of equation (13), set equal to zero, for \mathbf{X} , yielding:

$$\mathbf{X}^{k+1} = (\mathbf{H}^\top \mathbf{W} \mathbf{H} + (\mu_X + \rho_X) \mathbf{I})^{-1} (\mathbf{H}^\top \mathbf{W} \mathbf{L} + \rho_X \mathbf{Z} - \mathbf{Y}_X^k). \quad (16)$$

Due to the spectral dependence of \mathbf{W} , the minimization with respect to \mathbf{X} must be carried out spectral line by spectral line so that equation (16) is used to find \mathbf{X} by finding \mathbf{X}_λ using each unique \mathbf{W}_λ ; however, this step can be parallelized and the inversion step is constant. The pseudo-code for solving equation (16) is shown in Algorithm 2.

Algorithm 2: Pseudocode for TLDR's minimization of the Augmented Lagrangian with respect to \mathbf{X} .

```

for  $\lambda$  do
     $\mathbf{X}_\lambda =$ 
     $(\mathbf{H}^\top \mathbf{W}_\lambda \mathbf{H} + (\mu_X + \rho_X) \mathbf{I})^{-1} (\mathbf{H}^\top \mathbf{W}_\lambda \mathbf{L}_\lambda + \rho_X \mathbf{Z}_\lambda - \mathbf{Y}_{X,\lambda})$ 
end
    
```

5.3 Minimization w.r.t. regularizers

The minimization of the augmented Lagrangian for each of the regularized variables involving the ℓ_1 -norms of \mathbf{N} and \mathbf{T} is identical, though the input changes. The minimization problem can be written using a stand-in variable \mathbf{r}_i ,

$$\arg \min_{\mathbf{r}_i} f(\mathbf{r}_i) + \frac{\rho_{r_i}}{2} \|\mathbf{r}_i - \tilde{\mathbf{r}}_i\|_2^2, \quad (17)$$

when posing $\tilde{\mathbf{r}}_i = \mathbf{Z} - \mathbf{Y}_{r_i} / \rho_{r_i}$. This is as a MAP problem, solvable by proximal operators. A proximal operator allows the minimization to be replaced with analytic functions that can be evaluated quickly within the ADMM loop, e.g. the proximal operator for the ℓ_1 -norm is known as soft-thresholding (Donoho 1995). The solution to the minimization is the projection

$$\mathbf{r}_i^+ = \text{sgn}(\tilde{r}_i) \max\left(|\tilde{r}_i| - \frac{\mu_r}{\rho_r}, 0\right). \quad (18)$$

Though the projection is identical, the input for the projected solution for \mathbf{T} must include the gradient. So that for \mathbf{T} , $\tilde{\mathbf{r}}_T = \nabla \mathbf{Z} - \mathbf{Y}_T / \rho_T$. In practice, the projection relating to negative values $\tilde{r}_i \leq -\mu_r / \rho_r$ in the VDM are rejected as an additional enforcement of positivity.

Similarly, for the positivity regularizer, equation (17) holds and only the projection changes. Where the projection is simply

$$\mathbf{P}^+ = \max\left(0, \mathbf{Z} - \frac{\mathbf{Y}_P}{\rho_P}\right). \quad (19)$$

In effect, the proximal operator for positivity simply rejects values where $\mathbf{Z} < \frac{\mathbf{Y}_P}{\rho_P}$.

5.4 Minimization w.r.t. \mathbf{Z}

To find the analytical solution for \mathbf{Z} , the same procedure as used for the minimization of \mathbf{X} is used. Begin by taking the derivative of

equation (13) with respect to \mathbf{Z} , set equal to zero and solve for \mathbf{Z} . Yielding,

$$\begin{aligned} \mathbf{Z}^{k+1} &= ((\rho_X + \rho_P + \rho_N) \mathbf{I} + \rho_T \nabla^\top \nabla)^{-1} \\ &(\mathbf{Y}_X + \rho_X \mathbf{X} + \mathbf{Y}_P + \rho_P \mathbf{P} + \mathbf{Y}_N + \rho_N \mathbf{N} + \nabla^\top \mathbf{Y}_T + \rho_T \nabla^\top \mathbf{T}). \end{aligned} \quad (20)$$

The minimization w.r.t. \mathbf{Z} benefits from the lack of spectral dependence in the solution. In its implemented form, \mathbf{Z} is calculated very quickly as matrix multiplication with fast Fourier transforms to handle the gradients. Implementing the forward difference gradient operators as an implementation of the convolution theorem allows the \mathbf{Z} update step to be calculated by only multiplication. Note also, that the inversion used is a constant, meaning that the \mathbf{Z} update step is extremely fast.

5.5 Convergence testing

In TLDR, convergence testing is carried out as outlined in Boyd et al. (2011, section 3.3.1). At each iteration k , we define the primal residual \mathbf{r}^k , dual residual \mathbf{s}^k , and their tolerances ϵ^{pri} and ϵ^{dual} as

$$\mathbf{r}^k = \left(\mathbf{X}^k - \mathbf{Z}^k, \mathbf{N}^k - \mathbf{Z}^k, \mathbf{P}^k - \mathbf{Z}^k, \frac{1}{2} (\mathbf{T}^k - \nabla \mathbf{Z}^k) \right) \quad (21)$$

$$\mathbf{s}^k = (\rho_X (\mathbf{Z}^k - \mathbf{Z}^{k-1}), \rho_N (\mathbf{Z}^k - \mathbf{Z}^{k-1}), \rho_P (\mathbf{Z}^k - \mathbf{Z}^{k-1}), \rho_T (\mathbf{Z}^k - \mathbf{Z}^{k-1})) \quad (22)$$

$$\begin{aligned} \epsilon^{\text{pri}} &= \epsilon^{\text{abs}} + \epsilon^{\text{rel}} \max \{ \|\mathbf{X}^k\|_2, \|\mathbf{N}^k\|_2, \|\mathbf{P}^k\|_2, \\ &\frac{1}{2} \|\mathbf{T}^k\|_2, \|\mathbf{Z}^k\|_2 \}, \end{aligned} \quad (23)$$

and

$$\begin{aligned} \epsilon^{\text{dual}} &= \epsilon^{\text{abs}} + \epsilon^{\text{rel}} \max \{ \|\mathbf{Y}_X^k\|_2, \|\mathbf{Y}_N^k\|_2, \|\mathbf{Y}_P^k\|_2, \\ &\frac{1}{2} \|\nabla^\top \mathbf{Y}_T^k\|_2 \}, \end{aligned} \quad (24)$$

where ϵ^{pri} and ϵ^{dual} are the global tolerances on the primal and dual residuals respectively.

TLDR is said to have converged when both the norm of the primal residual and norm of the dual residual have dropped below their respective tolerances ϵ^{pri} and ϵ^{dual} . Which is to say, TLDR has converged when $\|\mathbf{r}^k\|_2 \leq \epsilon^{\text{pri}}$ and $\|\mathbf{s}^k\|_2 \leq \epsilon^{\text{dual}}$. These parameters can be tuned to get the desired convergence behaviour. Should the convergence criterion not be met, TLDR will terminate after maximum number of iterations set by the user.

5.6 Input data

Intended as a flexible method for reconstructing VDMs from RM data, TLDR makes minimal assumptions of the data provided to. As mentioned in Section 2, the continuum data must exist on an appropriate grid of times corresponding to the sample times of the spectral data and the delay times in the VDM to be reconstructed. In order to guarantee proper construction of the convolution operator TLDR applies a linear interpolation to the continuum data, but does not consider the uncertainty on the continuum. This is the only processing applied to continuum data by TLDR. Should continuum data be modelled at the appropriate times before being passed to TLDR, this linear interpolation will not affect the data in any way. Additionally, TLDR makes no assumptions on the spectral data provided to it. The option to model either the continuum data, or the spectra data is left to the user as data preparation.

While TLDR does not require modelled data to be provided to it, modelling of both continuum data and spectral data prior is strongly recommended prior to reconstructing a VDM. As TLDR does not consider uncertainties on continuum data, modelling the continuum allows these uncertainties and large gaps in observations to be considered before any VDM is reconstructed, as well as providing a reliable method for interpolating the light curve. Gaussian process modelling, previously adopted for regularized linear inversion (Skjelboe et al. 2015), and damped random walk models, as implemented in the JAVELIN code (Zu et al. 2011), have successfully been used for modelling continuum light curves for RM. For a fast Gaussian process modelling method easily applied to RM data, see Foreman-Mackey et al. (2017). Similarly, modelling of spectral data with software such as PrepSPEC can isolate and remove constant spectral components, allowing TLDR to reconstruct a VDM based on only the broad-line variations in the data. A recent discussion of the PrepSPEC software’s spectral decomposition method can be found in Horne et al. (2021).

5.7 Additional TLDR notes

In order to streamline the use of TLDR and ease frustrations that arise due to the sensitivity of the algorithm to the selection of the penalty parameters, three strategies for automatically adapting the penalty parameters were tested. Unfortunately, the simple strategy in Boyd et al. (2011, section 3.4.1), the more advanced residual balancing scheme described in Wohlberg (2017), and the adaptive consensus ADMM (ACADMM) strategy of Xu et al. (2017) resulted in instabilities in the algorithm. These instabilities arise out of the proximal projections used in TLDR, though they worked quite well when these terms were not used. As such, TLDR does not have a method of automatically selecting penalty parameters. This means that TLDR contains seven individual parameters that must be tuned by the user, which can be an arduous process for the user.

Tuning can be broken down into two categories: the penalty parameters ($\rho_X, \rho_P, \rho_N, \rho_T$) and regularization hyperparameters ($\mu_{\ell_2}, \mu_{\ell_1}, \mu_T$). Increasing the magnitude of the penalty parameters reduces the number of iterations required before convergence. Varying the ratio between the various penalty parameters changes the relative importance of the different regularizers. Changing the regularization hyperparameters increases the amount of smoothness or sparsity in the final VDM.

While tuning the various parameters for TLDR can be difficult, experience with the algorithm has revealed a general strategy. As a first step, set the penalty parameters to the same value and scale that value until TLDR converges in approximately 5000 iterations. Second, tune the regularization hyperparameters. Beginning with each regularization parameter set to 1.0, tuning should be carried out one regularizer at a time with quadratic smoothing μ_{ℓ_2} first, sparsity μ_{ℓ_1} second, and TV μ_T last. Increase the value of the regularization parameter until the final reduced- χ^2 reaches an inflection point and repeat for the next regularizer. Third, tune the relative ratio of the penalty parameters to move the final reduced- χ^2 towards unity. Generally, this step has required increasing the penalty parameter for positivity more than anything else. It is likely that the user will need to iterate on the second and third steps to achieve a satisfactory result.

The TLDR algorithm is written in the Julia programming language (Bezanson et al. 2017). For ease and functionality, regularization variables exist as data structures within the algorithm’s code. This allows the variables associated with each term to be stored and passed

in the algorithm as a part of the regularization variable to which they belong.

6 TESTING

In order to test TLDR, synthetic emission line data have been created using a number of model VDMs convolved with a continuum light curve. To generate the synthetic emission line data the Johnson B band continuum light curve for Arp 151 from the Lick AGN Monitoring Project (LAMP) is adopted, the complete discussion of the photometric observing campaign and data reduction can be found in Walsh et al. (2009). Observations of Arp 151 took place between 2008 February 10 and May 16 and the resulting data were calibrated via differential photometry. The original continuum data have further been reduced by binning any observations made within 0.1 d. The resulting light curve consists of 66 observations with a median sampling frequency of 1.02 d, mean sampling frequency of 1.5 d, mean flux of $\langle f \rangle = 0.93$ mJy with a standard deviation of 0.15 mJy, and flux ratio $f_{\max}/f_{\min} = 1.8$. The light curve has an excess variance of $F_{\text{var}} = 0.16$ calculated as

$$F_{\text{var}} = \frac{\sqrt{\sigma^2 - \delta^2}}{\langle f \rangle}, \quad (25)$$

where σ^2 is the variance of the flux and δ^2 is the mean square of the uncertainties of the flux. The continuum light curve is linearly interpolated on to a grid based on the date of spectroscopic observation and delay sample times. The dates used for the spectroscopic observations are those from the spectroscopic monitoring of Arp 151 from the LAMP campaign which range from 2008 March 25 to May 21 with a median sampling of 1.02 d and a mean sampling of 1.4 d, the full description of the spectroscopic observations and data reduction are provided in Bentz et al. (2009). For all tests a delay sampling time of 1.0 d was used for the reconstruction. From the interpolated continuum light curve the convolution operator is created. Synthetic emission line data is then generated and Gaussian noise is added,

$$\mathbf{L}_{\text{synthetic}} = \mathbf{H}\mathbf{X}_{\text{synthetic}} + \mathbf{n}. \quad (26)$$

The continuum data used to generate all synthetic emission line data used for testing are shown in the top plot in Fig. 5 where the black circles represent the continuum observations of Arp 151 with error bars included and the black line represents the linear interpolation of those observations. By feeding the synthetic emission line data and the real continuum data into TLDR, the synthetic VDM is reconstructed. This provides a measure of the reconstructive efficacy of TLDR. Using the continuum data and sampling dates from actual observations of Arp 151 in the creation of the artificial emission line data creates a data set with realistic cadences, but none of the following tests are intended to be compared to any reconstruction of the VDM for the Arp 151 data set.

6.1 Basic tests

In order to test the behaviour of the TLDR algorithm a number of simple VDMs were used. Each test VDM consists of a 10×10 pixel matrix with signal values ranging between a minimum of 0.0 and a maximum of 1.0. The test data sets were prepared with emission line errors yielding an SNR of 50.0. These test VDMs and their reconstructions shown in Fig. 2 where the original VDM appears on the left of each pair with the reconstruction on the right. On the far right is the colour map used for all of the test reconstructions, this colour map was specifically chosen to emphasize deviations

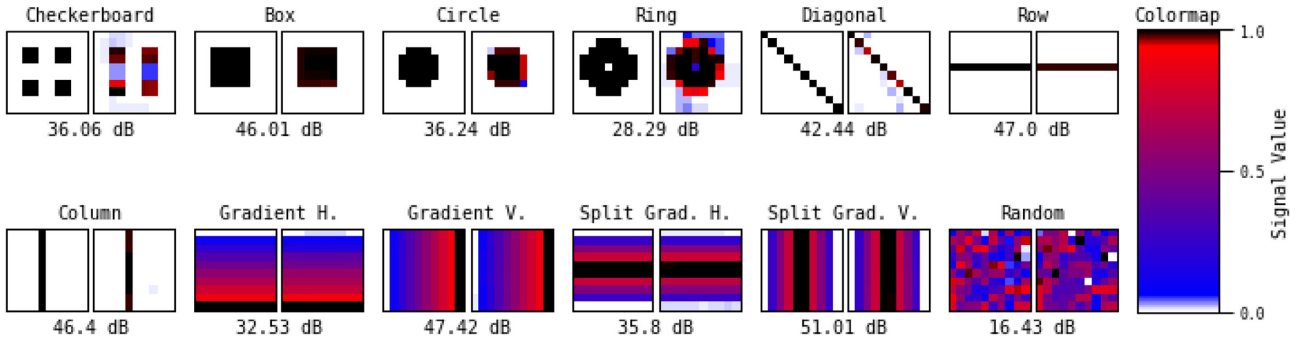


Figure 2. Reconstruction of 10×10 pixel test VDMs with signal values ranging from 0.0 to 1.0 and a signal-to-noise ratio of 50 in the synthetic emission lines. In each pair, the left plot shows the test VDM and the right plot shows the reconstructed VDM. Beneath each pair is the PSNR value of the reconstructed VDM in decibels. On the far right is the colour map used for each inset plot which emphasizes deviations from the extrema of the test VDMs making problem areas in the reconstructed VDMs stand out.

from the expected values in the VDMs. Below each pair the peak signal-to-noise ratio (PSNR) of the reconstructed VDM is displayed providing a measure of the reconstruction fidelity. Where the PSNR is calculated in the usual way with

$$\text{PSNR} = 20 \cdot \log_{10} \left(\frac{1}{\sqrt{\text{MSE}}} \right) \quad (27)$$

and

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N (\mathbf{X}_i - \mathbf{X}_{\text{true},i})^2. \quad (28)$$

The geometries were chosen to test the ability of TLDR to recover characteristics which may be encountered in real data. Clearly, the algorithm excels at reconstructions of smooth gradients and places where test VDMs have continuous values across their cross-sections. This is unsurprising as the majority of the regularizers used in TLDR promote smoothness. The algorithm has some difficulty when the VDM being reconstructed contains abrupt signal change as in the Checkerboard, Box, Circle, Ring, and Diagonal test VDMs. Unsurprisingly, TLDR struggles most when reconstructing a VDM consisting purely of random noise.

6.2 Synthetic VDM of Keplerian disc

In order to test TLDR with a more complicated VDM than those used in Section 6.1, a synthetic VDM was created using a simple flat Keplerian disc toy model of the BLR. This synthetic VDM was generated from an array of 1 million randomly situated points arranged in a flat disc about a central point corresponding to the location of an SMBH with mass $7 \times 10^6 M_{\odot}$ with an inner radius of 1.0 and an outer radius of 5.0 light-days.

As demonstrated in other work on RM (Welsh & Horne 1991; Peterson & Horne 2004), the line-of-sight velocities \mathbf{v} and delay times τ for particles in an edge-on circular disc $i = 90$ are easily found from the geometry and the orbital velocity V at radius r , where the angle ϕ is the azimuth angle in the plane of the disc starting from zero at the line of sight to the observer.

$$\mathbf{v}(r, \phi) = V(r) \sin(\phi) \quad (29)$$

and

$$\tau(r, \phi) = \frac{r}{c} (1 - \cos(\phi)). \quad (30)$$

After calculating the line-of-sight velocities and delays for all of the points in the simulated disc, the points are projected on to the VDM

seen in Fig. 3 middle panel. The VDM is then binned along both delay time and velocity (Fig. 3 lower panel) with the delay times scaled to a maximum of 50.0 light-days. The signal value of each pixel is normalized and scaled by counting the number of points that lie within a given pixel, normalizing to the maximum number of points found in a single pixel, and scaling the normalized array to the desired maximum value. Finally, each pixel is scaled by the ratio of its delay to the maximum delay to enhance the structure that appears in the VDM. In this case, the VDM was binned with 50 pixels of 1.0 light-day along the temporal axis and 20 pixels on the spectral axis which happen to be 597 km s^{-1} wide, with a maximum signal value of 0.1.

This toy model is oversimplified in that the binned flux does not account for the relative flux difference in different delays corresponding to the light-travel differences. Furthermore, the scaling and binning used for the reconstruction change the physicality of the model. Similarly, the use of a flat Keplerian disc is simplistic. The goal in generating this synthetic VDM is to provide a test case showing the algorithm's capability for recovering complicated signals that mimic those recovered by other methods. This synthetic VDM is not intended as a realistic model of the BLR. Furthermore, despite borrowing much from the Arp 151 data set to create this synthetic line emission data, this toy model is not intended to provide any physical comparison to any reconstruction of a VDM for the Arp 151 data set.

The synthetic VDM was used to create synthetic spectra according to equation (26) with a fractional uncertainty of 1.5 per cent matching the mean fractional error from the real Arp 151 data set (Bentz et al. 2009) and then reconstructed using TLDR. The reconstructed synthetic VDM generated from the Keplerian disc model and Arp 151 continuum is shown alongside the original synthetic VDM in Fig. 4. Comparing the reconstructed Keplerian disc VDM to the original shows that the algorithm is providing an excellent reconstruction. There is some noisy distortion apparent at longer delay times and some loss of fine detail near the edges of the VDM, but with a PSNR of 45.0 dB the details of the VDM are well reconstructed.

The reconstruction can be further evaluated by looking at a sample of the light curves from the reconstruction shown in Fig. 5. The top plot is simply the continuum light curve used to generate the synthetic emission line data (i.e. the original Arp 151 continuum data) shown by black circles with error bars slightly larger than the markers and the interpolated data shown as the black solid line. The bottom four rows show the spectral light curves from the data at 20 per cent (-1196 km s^{-1}), 40 per cent (-2394 km s^{-1}), 60 per cent




Figure 3. Top: Overview of the simulated flat Keplerian disc showing 10 000 particles in the same two-dimensional isodelay surface as Fig. 1. The size of the markers shows relative delay and colour indicates redshift/blueshift due to rotational velocity. Middle: Projected VDM of the particles in the simulation, again colour indicates redshift/blueshift. Bottom: Binned VDM of the simulated data.

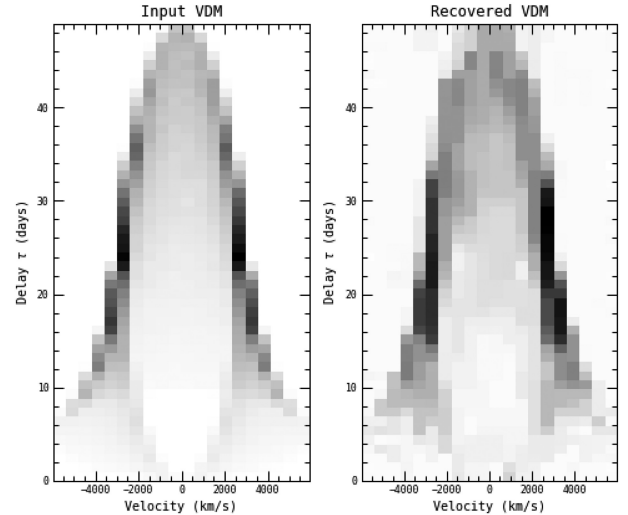


Figure 4. Synthetic VDM of a Keplerian disc (left) and the recovered VDM (right). Recovered VDM shows much of the detail of the synthetic VDM, with some artefacting visible in the recovered VDM.

(-3591 km s^{-1}), and 80 percent (-4789 km s^{-1}) of the maximum line-of-sight rotational velocity of the model. Each shows the response function on the left with the true response as the grey dashed line and the reconstructed response as the green solid line. On the right each shows the spectral light curve with the noiseless light curve shown as a black dashed line, the noisy spectral light curve input as the blue circles, and the recovered spectral light curve as the red solid line. Each also contains the line's reduced χ^2 value. While in each case, the recovered spectral line closely matches the noiseless true spectral line, they all show deviation from the true response function in the recovered response functions. The second row showing the curves for the -1196 km s^{-1} cross-section shows a very good recovery of the response function and spectral light curve. The bottom three rows, showing the -2394 , -3591 , and -4789 km s^{-1} lines, contain oversmoothed responses where the signals assume the general shape expected but are suppressed in the regions of maximum response value. It would likely be possible to achieve a better fit across all lines by using a different set of regularization hyperparameters for each line, but this would further complicate what can already be an arduous parameter selection process, a trade-off which must be considered for any image reconstruction process.

7 PRELIMINARY RECONSTRUCTION OF ARP 151 $H\beta$

As a final indication of the utility of TLDR, an initial reconstruction of the VDM for the $H\beta$ feature of Arp 151 from LAMP (Bentz et al. 2009) is provided. Using the continuum subtracted model spectra and the Johnson B continuum data used in for the maximum entropy method reconstruction presented in Bentz et al. (2010), an initial VDM was calculated in accordance with Section 5.1 with a smoothing hyperparameter of 2.0×10^6 . The VDM was set up with the pixel spectral width matching the 2.0 \AA steps of the original spectral observations and a temporal width of 1.0 d matching the global average sampling frequency of the continuum and spectral observations. For this reconstruction, TLDR was run for 3.0×10^4 iterations with the parameters listed in Table 2, selected from a grid of possible regularization hyperparameters by proximity to $\chi^2 = 1.0$.




Figure 5. Sample light curves from the reconstruction of the Keplerian disc. Top: Continuum light curve where the black circles show the original sample data with error bar and the black solid line shows the linear interpolation of the original data. Bottom four rows: On the left, response function where the grey dashed line is the true function and the green solid line is the recovered function. On the right: Noiseless spectral light curve shown by the black dashed line, the noisy spectral light curve shown as the blue circles, and the recovered spectral light curve as the red solid line. The lines shown represent 20 percent (-1196 km s^{-1}), 40 percent (-2394 km s^{-1}), 60 percent (-3591 km s^{-1}), and 80 percent (-4789 km s^{-1}) the maximum line-of-sight rotational velocity of the Keplerian disc model used to generate these emission line data.

Table 2. Reconstruction Parameters used in the reconstruction of the $H\beta$ feature of Arp 151.

Regularizer	Hyperparameter μ	Penalty parameter ρ
Smoothing X	2.0×10^6	1.0
Positivity P	N/A	100.0
Sparsity N	10.0	1.2
Total variation T	1.75	1.2

The reconstruction terminated with a reduced χ^2 of 1.02, yielding the VDM shown in Fig. 6. Visual comparison between the VDM reconstructed with TLDR and those recovered by the maximum entropy method (Bentz et al. 2010) and regularized linear inversion (Skjelboe et al. 2015) largely agree. In each case, the overall shape of the response in the VDM appears to match with asymmetries appearing in similar locations. However, to make any substantive comparison between results a thorough comparative analysis would be necessary.

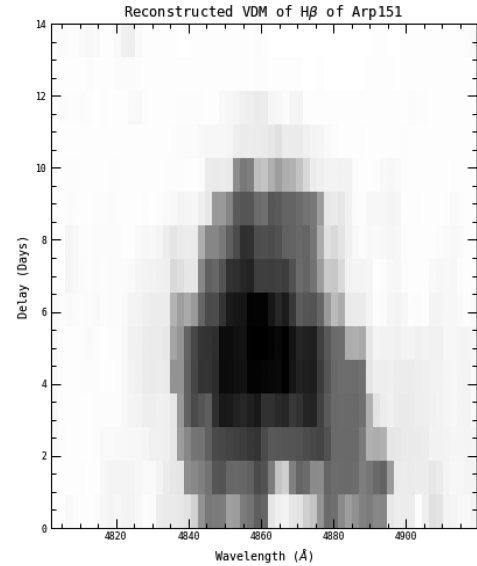


Figure 6. Reconstructed TDF of the $H\beta$ feature of Arp 151 from the LAMP RM project.

8 CONCLUSION

The TLDR algorithm provides an additional tool for RM, that can be used in concert with existing tools or independently, to reconstruct VDMs from RM data. The regularization scheme implemented within TLDR provides a flexible platform for reconstruction of VDMs. With numerous recent and ongoing RM campaigns (De Rosa et al. 2015; Shen et al. 2016; LAMP2016 Collaboration 2017), interest in using RM to study AGNs is high. With high interest, and a growing number of data sets, TLDR should be able to make a sizable contribution to the RM field.

Future work on TLDR will include applying the algorithm to a number of the well-studied data sets from the LAMP (Bentz et al. 2009). The data contained in the LAMP data release have been studied by all of the RM algorithms currently in use, including the maximum entropy method (Bentz et al. 2010), regularized linear inversion (Skjelboe et al. 2015), as well as direct modelling (Pancoast et al. 2014). This will provide a way to directly evaluate the performance of TLDR and compare the different methods currently available.

ACKNOWLEDGEMENTS

FB and MA acknowledge support from the National Science Foundation through AAG grants AST-1210972, AST-1616483, and AST-1814777. MCB gratefully acknowledges support from the National Science Foundation through CAREER grant AST-1253702 and AAG grant AST-2009230 to Georgia State University. We thank the referee for their thoughtful comments and thorough review.

DATA AVAILABILITY

The TLDR software repository is publicly available at <https://github.com/matdander/TLDR>. The data underlying this article were accessed from the Lick AGN Monitoring Project (LAMP2016 Collaboration 2017) at <https://www.physics.uci.edu/barth/lamp.html>.

REFERENCES

- Bentz M., 2009, Improving the Radius-Luminosity Relationship for Broad-Lined AGNs with a New Reverberation Sample. HST Proposal
- Bentz M. C., Katz S., 2015, *PASP*, 127, 67
- Bentz M. C. et al., 2009, *ApJ*, 705, 199
- Bentz M. C. et al., 2010, *ApJ*, 720, L46
- Bentz M. C. et al., 2013, *ApJ*, 767, 149
- Bezanson J., Edelman A., Karpinski S., Shah V. B., 2017, *SIAM Rev.*, 59, 65
- Blandford R. D., McKee C. F., 1982, *ApJ*, 255, 419
- Boley D., 2013, *SIAM J. Optim.*, 23, 2183
- Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
- Boyd S., Parikh N., Chu E., Peleato B., Eckstein J., 2011, *Found. Trends Mach. Learn.*, 3, 1
- Bryan K., Leise T., 2013, *SIAM Rev.*, 55, 547
- Candes E. J., Wakin M. B., 2008, *IEEE Signal Process. Mag.*, 25, 21
- Cao N., Nehorai A., Jacob M., 2007, *Opt. Express*, 15, 13695
- Cromwell R., Weymann R., 1970, *ApJ*, 159, 379
- Czerny B., 2006, in Gaskell C. M., McHardy I. M., Peterson B. M., Sergeev S. G., eds, ASP Conf. Ser. Vol. 360, AGN Variability from X-Rays to Radio Waves. Astron. Soc. Pac., San Francisco, p. 265
- Davenport M. A., Duarte M. F., Eldar Y., Kutyniok G., 2011, in *Compressed Sensing: Theory and Applications*. Cambridge Univ. Press, Cambridge
- De Rosa G. et al., 2015, *ApJ*, 806, 128
- Donoho D. L., 1995, *IEEE Trans. Inf. Theory*, 41, 613
- Eckstein J., 2012, RUTCOR Research Reports, Vol. 32. Rutgers University
- Ferrari A., Deguignet J., Ferrari C., Mary D., Schutz A., Smirnov O., 2015, in *SKA Pathfinder Radio Continuum Surveys (SPARCS)*
- Foreman-Mackey D., Agol E., Ambikasaran S., Angus R., 2017, *AJ*, 154, 220
- Gravity Collaboration, 2020, *A&A*, 635, A92
- Grier C. J., Peterson B. M., Horne K., Bentz M. C., Pogge R. W., Denney K. D. et al., 2013, *ApJ*, 764, 47
- Han D., Yuan X., 2013, *SIAM J. Numer. Anal.*, 51, 3446
- Hoerl A. E., Kennard R. W., 1970, *Technometrics*, 12, 55
- Horne K., 1994, in Gondhalekar P. M., Horne K., Peterson B. M., eds, ASP Conf. Ser. Vol. 69, Reverberation Mapping of the Broad-Line Region in Active Galactic Nuclei. Astron. Soc. Pac., San Francisco, p. 23
- Horne K., Peterson B. M., Collier S. J., Netzer H., 2004, *PASP*, 116, 465
- Horne K. et al., 2021, *ApJ*, 907, 76
- Kelly B. C., Bechtold J., Siemiginowska A., 2009, *ApJ*, 698, 895
- Krolik J. H., 1994, in Gondhalekar P. M., Horne K., Peterson B. M., eds, ASP Conf. Ser. Vol. 69, Reverberation Mapping of the Broad-Line Region in Active Galactic Nuclei. Astron. Soc. Pac., San Francisco, p. 53
- LAMP2016 Collaboration, 2017, in Vivian U., ed., American Astronomical Society Meeting Abstracts #229, p. 414.04
- Li M., Zhang C., Peng C., Guan Y., Xu P., Sun M., Zheng J., 2016, *BioMed Res. Int.*, 2016, 12
- Nishihara R., Lessard L., Recht B., Packard A., Jordan M. I., 2015, in Proceedings of the 32nd International Conference on Machine Learning, Vol. 37, ICML'15. JMLR.org, p. 343, <http://dl.acm.org/citation.cfm?id=3045118.3045156>
- Pancoast A., Brewer B. J., Treu T., 2011, *ApJ*, 730, 139
- Pancoast A., Brewer B. J., Treu T., 2014, *MNRAS*, 445, 3055
- Penston M. V., Penston M. J., Selmes R. A., Becklin E. E., Neugebauer G., 1974, *MNRAS*, 169, 357
- Peterson B. M., 1988, *PASP*, 100, 18
- Peterson B. M., 1993, *PASP*, 105, 247
- Peterson B. M., 1994, in Gondhalekar P. M., Horne K., Peterson B. M., eds, ASP Conf. Ser. Vol. 69, Reverberation Mapping of the Broad-Line Region in Active Galactic Nuclei. Astron. Soc. Pac. San Francisco, p. 1
- Peterson B. M., Horne K., 2004, *Astron. Nachr.*, 325, 248
- Peterson B. M., Wandel A., 2000, *ApJ*, 540, L13
- Rudin L. I., Osher S., Fatemi E., 1992, *Phys. D*, 60, 259
- Sandage A., 1967, *ApJ*, 150, L177
- Schutz A., Ferrari A., Mary D., Soulez F., Thiébaud É., Vannier M., 2014, *J. Opt. Soc. Am. A*, 31, 2334
- Schutz A., Ferrari A., Mary D., Thiébaud É., Soulez F., 2015, in EU-SIPCO NICE, France, <https://hal.archives-ouvertes.fr/hal-01123500>
- Shen Y. et al., 2015, *ApJS*, 216, 4
- Shen Y. et al., 2016, *ApJ*, 818, 30
- Skjelboe A., Pancoast A., Treu T., Park D., Barth A. J., Bentz M. C., 2015, *MNRAS*, 454, 144
- Storchi-Bergmann T., Schnorr-Müller A., 2019, *Nat. Astron.*, 3, 48
- Thiébaud E., 2006, in Foy R., Foy F. C., eds, *Optics in Astrophysics*, NATO Science Series II: Mathematics, Physics and Chemistry, 198. Springer, Dordrecht, p. 397
- Thiébaud E., Giovannelli J.-F., 2010, *IEEE Signal Process. Mag.*, 27, 97
- Thiébaud É., Soulez F., 2012, In *Optical and Infrared Interferometry III*. Proceedings of the SPIE, Volume 8445. SPIE, p. 450
- Tikhonov A. N., 1963, *Sov. Math. Dokl.*, 4, 1035
- Ulrich M.-H., Maraschi L., Urry C. M., 1997, *ARAA*, 35, 445
- Walsh J. L. et al., 2009, *ApJS*, 185, 156
- Wang Y., Yin W., Zeng J., 2018, *J. Sci. Comput.*, 78, 29
- Welsh W. F., Horne K., 1991, *ApJ*, 379, 586
- Wohlberg B., 2017, Admm Penalty Parameter Selection by Residual Balancing. preprint ([arXiv:1704.06209](https://arxiv.org/abs/1704.06209))
- Xu Z., Taylor G., Li H., Figueiredo M. A. T., Yuan X., Goldstein T., 2017, in Precup D., Teh Y. W., eds, Proceedings of the 34th International Conference on Machine Learning, Vol. 70. PMLR, p. 3841
- Yao J., Tian F., Rakkongthai Y., Oraintara S., Liu H., 2015, *Biomed. Opt. Express*, 6, 2961
- Zu Y., Kochanek C. S., Peterson B. M., 2011, *ApJ*, 735, 80

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.