

---

AI translation · View original & related papers at  
[chinaxiv.org/items/chinaxiv-202505.00087](https://chinaxiv.org/items/chinaxiv-202505.00087)

---

## Constraining Schwarzschild Models with Orbit Classifications (Postprint)

**Authors:** Richard J. Long

**Date:** 2025-04-28T11:58:00+00:00

### Abstract

A simple orbit classification constraint extension to stellar dynamical modeling using Schwarzschild's method is demonstrated. The classification scheme used is the existing "orbit circularity" scheme ( $\lambda z$ ) where orbits are split into four groups—hot, warm, cold, and counter-rotating orbits. Other schemes which can be related to the orbit weights are expected to be viable as well. The results show that the classification constraint works well in modeling. However, given that orbits in external galaxies are not observable, it is not clear how the orbit classification for any particular galaxy may be determined. Perhaps range constraints for different types of galaxies determined from cosmological simulations may offer a way forward.

### Full Text

#### Preamble

Research in Astronomy and Astrophysics, 25:035011 (12pp), 2025 March © 2025. National Astronomical Observatories, CAS and IOP Publishing Ltd. All rights, including for text and data mining, AI training, and similar technologies, are reserved. Printed in China. <https://doi.org/10.1088/1674-4527/ad9bfe> CSTR: 32081.14.RAA.ad9bfe

#### Constraining Schwarzschild Models with Orbit Classifications

Richard J. Long

Department of Astronomy, Tsinghua University, Beijing 100084, China; rjlas-tro@yahoo.com

Received 2024 October 19; revised 2024 December 1; accepted 2024 December 6; published 2025 March 7

## Abstract

A simple orbit classification constraint extension to stellar dynamical modeling using Schwarzschild’s method is demonstrated. The classification scheme used is the existing “orbit circularity” scheme ( $\lambda z$ ) where orbits are split into four groups—hot, warm, cold, and counter-rotating orbits. Other schemes which can be related to the orbit weights are expected to be viable as well. The results show that the classification constraint works well in modeling. However, given that orbits in external galaxies are not observable, it is not clear how the orbit classification for any particular galaxy may be determined. Perhaps range constraints for different types of galaxies determined from cosmological simulations may offer a way forward.

**Key words:** galaxies: kinematics and dynamics – galaxies: general – methods: numerical

## 1. Introduction

The stellar dynamics of galaxies may be investigated by creating models of galaxies, constraining them with observational data, and then examining the models to see what might be learnt about the real galaxies. The book by Binney & Tremaine (2008) describes the theory behind some of the traditional modeling techniques that might be used. A key point which must be accepted about external galaxy models is that the models can only be indicative or illustrative of the galaxies because of the current technical limitations of the instruments employed for observations: six-dimensional (6D) galaxy phase space data are not obtainable. Mathematically, recovery of the phase space data from observations is not possible in general.

Related deprojection issues are included, for example, in Rybicki (1987), and more recently in papers by Cappellari (2020) and Vasiliev & Valluri (2020). Long et al. (2021) demonstrated that Schwarzschild’s method does not fully comply with the theoretical statements in Rybicki (1987). Observing and modeling our Galaxy are different in that the instrumental capability to obtain three-dimensional (3D) spatial and velocity data is available. As a consequence, for example, it is not clear to what extent orbits in galaxies in cosmological simulations are truly representative of those present in real galaxies. Specialized n-body simulations help us gain insight into the orbits that might comprise particular components of galaxies, e.g. galactic bars, but do not give us the capability to model all aspects of real galaxies accurately.

It has been recognized that applying machine learning modeling methods to physical systems requires that relevant physical laws should be adhered to, and this has resulted in the development of physics informed neural networks (Raissi et al. 2019). It is appropriate to revisit more traditional stellar dynamical modeling methods, such as Schwarzschild’s method (Schwarzschild 1979) or Syer & Tremaine’s made-to-measure method (Syer & Tremaine 1996), to understand whether their solutions are astrophysically realistic or not. In this investigation,

Schwarzschild's method which superimposes weighted stellar orbits to try and reproduce observational data is examined.

Using Schwarzschild's method to estimate the global properties of galaxies (e.g. total mass) is well-developed, but working with the weighted orbits comprising the Schwarzschild solution, perhaps to help understand the evolutionary history of a galaxy, does need careful consideration (Jin et al. 2019, 2020). The process used to determine the orbit weights is dealing with an under-determined optimization problem: there are few observational constraints by comparison with the number of orbits. Also, it must be remembered that Schwarzschild's method was created to demonstrate that some self-consistent models of triaxial galaxies could be created: no claims were made on how astrophysically realistic such models would be (Schwarzschild 1979; Binney & Tremaine 2008).

Operationally, the known shortcomings of Schwarzschild's method include the following:

1. **Low orbit utilization (without regularization)** in that only 10% of orbits supplied are an effective part of any solution—see Table 2 of Jin et al. (2019). This percentage becomes even lower if chemo-dynamical techniques such as in Long & Mao (2018) are applied. Regularization can be used to improve the utilization to typically 75% or more depending on the optimization method.
2. **The orbit weighting scheme is purely a numerical process** with the resulting weights showing a large weight imbalance in that a small number of orbits receive high weights and conversely many orbits have low or zero weights (Long & Mao 2018).
3. **Significant overfitting of galaxy data** with mean  $\chi^2$  values of  $10^{-1}$  or less is not uncommon.
4. **Difficulties in determining confidence intervals on model parameters** (Binney & Tremaine 2008, section 4.7.2).

In short, the orbit weights have numerical importance to the optimization process (weights act to scale individual orbit contributions to model observables) but whether the orbits individually or collectively as a set are realistic physically and truly exhibit the orbit structure of the galaxy being modeled remains unclear. There are no “necessary and sufficient” statements on representation in the literature regarding the orbits selected. The weighted orbits can only reflect what is required by the optimization's objective function: if the function is concerned with reproducing observables then the orbits are important in that context but that is all. As noted in Jin et al. (2019), for example, they are illustrative or indicative of what a galaxy might be like but cannot be taken as definitive or complete. By extrapolation, any orbit classification applied to the galaxy has to be caveated in the same way.

Given that Schwarzschild's method cannot be used to determine the overall orbit classification for a galaxy without amendment or enhancement, if the orbit

classification were to be known from other sources, could that classification be included in Schwarzschild modeling as a constraint perhaps? Answering this question is the focus of this article. The article's structure is as follows. Section 2 describes at a top level the approach taken in the investigations. The galaxy data used are described in Section 3, and the relevant theory and descriptions of the methods are in Section 4. Results and subsequent discussion are in Sections 5 and 6, with conclusions in Section 7.

## 2. Approach

Whereas the Introduction, Section 1, sets the context and describes what will be investigated, this section describes how it will be done. A number of axisymmetric models of galaxies are constructed, under varying conditions, using an existing software implementation of Schwarzschild's method taken from Long et al. (2021), and analyzed. Line-of-sight luminosity and kinematical data from nine galaxies in total are used, with four real galaxies from ATLAS<sup>3</sup>D (Capellari et al. 2011) and five simulated galaxies from IllustrisTNG (Nelson et al. 2019). The simulated galaxies have known orbit circularity classifications (Xu et al. 2019).

The initial conditions (spatial and velocity) for the orbits are constructed as in Long et al. (2021). The three integral (3I) scheme is used with simple adjustments to reflect the degree of bulk rotation of the galaxies, or to vary the number of circular orbits. Orbit weights are determined by employing Lawson & Hanson's non-negative least squares method which uses an active set algorithm internally (NNLS). The objective function is extended to include not only the usual reproduction of the target observational data but classification as well. Other optimization schemes could have been used but the NNLS scheme (without regularization) is computationally fast enough to be used in a grid search for orbit classification values.

Modeling runs are accomplished in three stages: 1) orbit integration and the collection of orbit contributions to the model observables, 2) weight determination using the NNLS optimization method, and 3) analysis of the modeling run, including the calculation of observable <sup>2</sup> values and examining the degree of conformance to the target orbit classification. In view of the under-determined nature of the problem, regularization may be used to restrict the weight solution space. Here, some models are run both with and without regularization so the impact of regularization is understood. Regularization takes two forms. The first penalizes heavy weights, and acts to increase the number of orbits actively contributing to the objective function. The second smooths the particle weights in integral space by penalizing high second derivative values. These two forms of regularization can be used separately or combined to give a composite form.

There are aspects of stellar dynamical modeling with Schwarzschild's method that are out of scope for investigation, for example chemo-dynamical modeling (Long & Mao 2018) or the use of discrete data (Chanamé et al. 2008).

Note that only axisymmetric (oblate) galaxy models are considered: the orbit weighting numerical optimization processes are not affected by choice of gravitational potential. Also, orbit classifications in this work are not targeted at addressing any issues to do with self-consistency, or data deprojection non-uniqueness.

### 3. Galaxy Data

For the ATLAS<sup>3</sup>D galaxies (NGC 1248, NGC 3838, NGC 4452, and NGC 4551), the data have been produced by the ATLAS<sup>3</sup>D survey (Cappellari et al. 2011). The four galaxies are the same galaxies as used in Long (2016) and Long & Mao (2018), and have various inclinations, different mass-to-light ratios, and sense of bulk rotation (see Table 1). Surface brightness values are calculated, using multi-Gaussian expansions (MGEs) (Section 4.1), on a (16, 16) polar grid with equal interval radial distances, and no straddling of a galaxy's major axis. The integral field unit (IFU) cells for kinematical data (first and second velocity moments) are the Voronoi cells resulting from signal-to-noise processing of the raw data.

The same IllustrisTNG galaxies (Nelson et al. 2019) as in Long et al. (2021) are used. The selection strategy for the galaxies is described in Long et al. (2021), and is not repeated here. The five axisymmetric galaxies (with identifiers, A490, A1090, A1190, A1290, and A1390) are viewed edge on with the mass-to-light ratio taken as 1.0 for all galaxies. It has been assumed that the kinematics of stellar particles can validly be used to create stellar kinematical data as if from a real galaxy (consistent with Li et al. (2016) or Xu et al. (2019), for example). Surface brightness values are calculated, using MGEs, on a (16, 16) polar grid with logarithmic interval radial distances, and straddling the galaxy's major axis. The IFU cells for kinematical data (the Gauss-Hermite coefficients  $h_1$  to  $h_4$ ) are constructed to mimic a real galaxy's data. Table 2 is the equivalent of Table 1 for the simulated galaxies.

The IFU cells for both the ATLAS<sup>3</sup>D galaxies and the IllustrisTNG simulated galaxies contain the observational kinematical data that Schwarzschild's method is asked to replicate. Data are symmetrized appropriately for an axisymmetric model/potential. The design matrix size used in the convex optimization (see equation 1) may be calculated as the total number of data points for observables times the number of orbits. As an example, NGC 3838 has surface brightness (256 points) and two IFU observables ( $2 \times 383 = 766$  points) giving 1022 data points in total. For 8000 orbits, the design matrix size is 1022 rows and 8000 columns. For simulated galaxy A1190 with four IFU observables, the matrix size is 696 rows and 8000 columns.

#### 4.1. Gravitational Potentials

Galaxy gravitational potentials and derivatives are derived from MGEs of the galaxies' surface brightness (Emsellem et al. 1994; Cappellari 2002). For com-

puter performance reasons, an MGE potential and its gradients are implemented as a set of interpolation tables, for example as in van den Bosch et al. (2004) or Long & Mao (2012).

For the simulated galaxies, dark matter potentials (calculated from the simulation's dark matter particles) are included as spherical Gaussians. No black hole modeling takes place. For the ATLAS<sup>3D</sup> galaxies, no dark matter is included (for consistency with Cappellari et al. (2013)), and central black holes are modeled as point sources.

#### 4.2. Orbit Initial Conditions

The initial phase space coordinates for orbits are created using the 3I scheme. Three variants are created by varying the sense of rotation of the orbits: 1) no overall sense of rotation about the symmetry axis, that is 50% of orbits with one sense of rotation and 50% with the other; 2) with the sense of rotation matching that shown by a galaxy's mean line-of-sight velocity data for 75% of the orbits with 25% counter-rotating; and 3) as in (2) but with 100% matching the data's sense of rotation and no counter-rotating orbits.

Similarly to Zhu et al. (2018), the circularity measure  $\lambda_z$ , the ratio of the orbit's  $L_z$  component of angular momentum to the maximum value allowed by the orbit's energy, is used to classify orbits into one of hot ( $\lambda_z \leq 0.25$ ), warm ( $0.25 < \lambda_z < 0.8$ ), cold ( $\lambda_z \geq 0.80$ ), or counter-rotating (the remainder). Practically, internally, it is more convenient to split the orbits into six classes—rotating or counter-rotating and then one of hot, warm or cold—and only combine classes where necessary for presentational or comparison purposes. These classes are used extensively in analyzing modeling runs.

The total orbit weights in each of the classes are taken as the overall orbit classification for a galaxy.

Note that the combination of a steady state gravitational potential for a galaxy and the initial spatial and velocity conditions for an orbit means that the orbit's trajectory is fully determined. From Hamiltonian mechanics, energy is a conserved quantity for an orbit, and, in addition, for an axisymmetric galaxy potential so is the component of angular momentum ( $L_z$ ) parallel to the symmetry axis of the galaxy. This means that the  $\lambda_z$  classifier of orbits as hot, warm, or cold can be found directly from the initial conditions. As a consequence, if all the orbits in a model are regarded as active orbits (not zero-weighted) by the optimization method then the relative proportions of the numbers of hot, warm, and cold orbits must match the proportions from the initial conditions. This remains true even for gravitational systems where  $L_z$  is not conserved but some time or orbit averaged value is used instead. What is of more interest is the relative proportions of the orbit weights allocated to hot, warm, and cold orbits as these proportions can be related to the global properties of a galaxy.

### 4.3. Schwarzschild Models

As previously stated, the implementation of Schwarzschild’s method is as in Long et al. (2021) and the details are not repeated here. For simulated galaxies, the observational data utilized are surface brightness with kinematics (Gauss-Hermite coefficients) as in Rix et al. (1997) and Cretton et al. (1999). For the ATLAS<sup>3</sup>D galaxies, modeling is with surface brightness and the first and second velocity moments. All models use 8000 orbits. No use is made of orbit dithering or point symmetries in collecting orbit contributions to model observables. Note that observational data are not available to be used to constrain individual orbits. If they were, they could be included in the objective function to be optimized. Mixes of orbits by type distribution can be included, as described below in Section 4.4, provided such techniques are linear superposition in the same way that other observables are. This is believed to be a new extension to Schwarzschild’s method, and so has its own section.

Orbit weights are determined using NNLS. Taking equation 1 in Long & Mao (2018), in matrix form, the expression to be minimized is  $\|Dw - K_2\|_2$ , where  $D$  is the “design” matrix giving the orbit contributions to the model observables,  $w$  represents the orbit weights to be determined, and  $K$  contains target observables. The L2-norm (Euclidean norm) is indicated by  $\| \cdot \|_2$ . Note that Schwarzschild’s method can only weight the orbits it has been provided with. In that respect, the gravitational potential and the orbit initial conditions play a substantial role in determining the eventual solution from Schwarzschild’s method.

### 4.4. Constraints using Orbit Classifications

In this section, the weighting process is first considered in optimization objective function terms (rather than matrices). Usually, the Schwarzschild objective function is framed in terms of reproducing the observed luminosity and kinematical data so the optimizing process selects and weights orbits to try and cause this reproduction to happen. If the objective function is framed to achieve a different end effect (for example, related to the distribution of some other quantity), a set of orbits appropriate to that objective would be selected and weighted. What is needed here is that not only are observables reproduced but that some astrophysically realistic mix of orbit types is used in so doing.

The circularity measure  $\lambda_z$  can be used to define orbit types and is a convenient place to start. Mathematically, an expression to be optimized might be  $\sum_i (\delta(i, C)w_i - T)^2$ , where  $C$  is an orbit type within the classification scheme being used, and  $T$  is the target weight to be met by the orbits in  $C$ . The selection function  $\delta(i, C)$  takes the value 1 if orbit  $i$  has type  $C$  and 0 otherwise. If there are uncertainties  $\sigma$  associated with the  $T$  then this can be accommodated simply as in expression 3. Structurally, expression (3) has the same  $\chi^2$  form as the observable data in the objective function.

In this article, expression (2) is used, together with an orbit classification parameter  $\alpha_c$ , to set the overall importance of the expression amongst all the other

terms in the function. Here,  $\alpha = 10^3$ .

Reverting now to matrix terms, each class of orbit type  $C$  is represented in matrix  $D$  (see equation 1) by a row with each column entry having value 1 if orbit  $i$  is in  $C$  and 0 otherwise. The corresponding entries in matrix  $K$  are the  $T$  values. In the context of  $\lambda z$  circularity classifications, the [hot, warm, cold, counter-rotating] target values might, for example, be [30%, 45%, 15%, 10%] with four matrix rows being required, and the entries per row indicating whether the individual orbits are of the associated type (entry value = 1) or not (entry value = 0). Note that the mechanism described is not specific to  $\lambda z$  classifications but could be used with other classification schemes as well. Also, the approach described here is readily adaptable to the made-to-measure method.

#### 4.5. Regularization

Two forms of regularization constraint (Tikhonov 1963) are employed, penalizing heavy weights and smoothing the weight distribution. Penalizing heavy weights is as described in Long & Mao (2018) and has two effects. It acts to reduce the maximum orbit weight, and as a consequence causes more orbits to have non-zero weights. Including this regularization term in equation 1 gives  $Dw - K_2 + \alpha w_2$ , where  $\alpha$  is a positive parameter controlling the amount of regularization. As a constraint,  $w \geq 0$ . Extending the regularization term to include priors  $W$  on the orbit weights changes the constraint term to  $w/W_2 = 1$ . Such regularization has been used elsewhere, for example, in Valluri et al. (2004) and Vasiliev (2013).

Smoothing the weight distribution is performed in integral space and is described below. For axisymmetric models, because of uncertainties about quite what form the third integral takes, only the conserved quantities energy ( $E$ ) and the  $Lz$  component of angular momentum are used. Working with the initial conditions on the orbits, each orbit's nearest two neighbors in  $(E, Lz)$  space are determined, and the second derivatives of the weights using finite difference formulae are calculated at the orbit's  $(E, Lz)$  position. A two-dimensional tree is used to find the nearest neighbors and the finite difference formula used (backward, forward or central) depends on the relative positioning of the orbit and its neighbors. Solution or curvature smoothing is a well-established technique, see Boyd & Vandenberghe (2004) or Press et al. (2007), for example.

The regularization parameters are  $1.6 \times 10^3$  for penalization and  $6.4 \times 10^2$  for smoothing, and were determined by experimentation. For convenience, the same values are used for all galaxies. Note that both forms of regularization can be used individually in isolation or combined together to give a composite regularization term. For the remainder of this article, the terms “regularization” (for heavy weights penalization) and “smoothing” (for second derivative smoothing) are employed.

## 5.1. Baseline Models

A set of models is constructed and analyzed to establish a baseline of understanding, or a reference point, for subsequent modeling. Both sets of galaxies are used for this task. 3I initial conditions, with the rotation of 75% of the orbits aligned to the bulk rotation of the galaxies, are employed. The orbit classification resulting from the weighted orbits is calculated and, for the simulated galaxies, compared with the classification in Xu et al. (2019).

The results for the simulated galaxies are shown in Table 3. Looking at the mean  $\chi^2$  column, the observed data values are reproduced well by the modeling. For models without regularization, mean  $\chi^2$  values  $> 1$  indicate that significant overfitting is taking place. The heaviest weighted orbits for a galaxy are those orbits remaining close (in projection) to the galaxy center with the lighter weighted orbits having larger apocenters. Figure 1 [Figure 1: see original paper] shows this weight progression for galaxy A1190 as an example. The position of the heaviest weights centrally matches the region where the surface brightness observable dominates numerically over other observables in the model. Using a simple circularity classification into hot, warm and cold orbits with no separation by sense of rotation, it is clear that orbits in different classes tend to have a role to play in particular regions—see Figure 2 [Figure 2: see original paper], this time using galaxy A1390 as an example.

As regularization is penalizing heavy weights, the maximum orbit weights are reduced (by approximately an order of magnitude), and the weight ranges are extended to lower values. Regularization does cause more orbits to be active. However, the issue concerning a small number of orbits being allocated heavy weights remains: it has just been “down-shifted” by a factor of 10—see the top row of Figure 3 [Figure 3: see original paper]. Smoothing alone also causes more orbits to be active but some orbits have even heavier weights as the orbit count distribution is broadened overall, generally to lower weights—see the bottom row of Figure 3. Combining regularization and smoothing together produces a result more akin to regularization alone but with more active orbits, and a performance penalty (see Table 8).

Overall, for the simulated galaxies, the models do not reproduce consistently the target orbit classifications from Xu et al. (2019). Also, the regularized and unregularized models themselves do not agree on their orbit classifications for a galaxy. In all cases, if the Xu et al. (2019) classification is used as a model observable, the model is able to reproduce that observable, demonstrating that the classification is not unachievable in a modeling sense. There is a minor difference regarding the determination of orbit classifications between this work and Xu et al. (2019): this work uses data extending to three effective radii while Xu et al. (2019) uses two effective radii. This difference has been assessed as not significant.

The results for the ATLAS<sup>3D</sup> galaxies are in Table 4, and are consistent with the simulated galaxy results except that there are no target orbit classifications

for comparison.

## 5.2. Orbit Classifications from Grid-searching

When no target orbit classification is available, it is possible to construct a grid search of Schwarzschild models to determine an appropriate classification from the orbit data. The search used here is to construct classifications by randomly uniformly selecting values from supplied minimum to maximum class ranges, and then to score those classifications using unregularized Schwarzschild modeling. A number of classifications closest in (Euclidean) distance to the classification with the minimum score are used to set revised minimum to maximum class ranges, and the classification construction and scoring processes are repeated. For the simulated galaxies, two iterations after the initial scoring are sufficient to have converged to an orbit classification for a galaxy.

The number of classifications created for scoring is 512, and 64 classifications are used to create the revised minimum to maximum class ranges. The [hot, warm, cold, counter-rotating] class ranges used initially for the simulated galaxies are [(0.2, 0.5), (0.2, 0.5), (0.1, 0.25), (0.1, 0.25)]. For the ATLAS<sup>3D</sup> galaxies, the range (0.1, 0.5) is used for all classes. Scores are calculated as the total  $\chi^2$  for the observables plus the orbit classification expression (2).

The results achieved are shown in Table 5. The search classifications match the classifications from the baseline unregularized modeling runs in Tables 3 and 6 quite well. This is to be expected from numerical optimization considerations. Pair-wise “corner” plots show a negative correlation between classification class values. This is not unexpected as the sum of the weights in the classes must be equal to 1.00, so an increase in one class must be compensated for in the other classes.

## 5.3. Initial Conditions—Rotation Extremes

In Section 5.1, all the models have initial conditions where there is a sense of bulk rotation aligned with the mean velocity data. In Table 6, a set of results for ATLAS<sup>3D</sup> galaxy models is shown where the models are lacking an overall sense of rotation. The actual ATLAS<sup>3D</sup> galaxies do have a sense of rotation—see Table 1. No sense of rotation implies that 50% of the orbits are rotating in the same sense as the velocity data and 50% are counter-rotating with the opposite sense. Within the table, counter-rotating orbits are split between the counter-rotation class (40%) and the hot class (10%). All the models appear to be acceptable (mean  $\chi^2 < 1$ ). The orbit classifications between Tables 4 and 6 do not agree, however. Our assessment is that, while it may be acceptable to allow the modeling method to decide what rotation mix is satisfactory for determining global galaxy attributes, for orbit classification investigations it is better to explicitly determine the sense of rotation which matches the observed data.

The fully aligned scheme, where all orbits are aligned with a galaxy’s sense of rotation, does not perform well: all simulated galaxy models fail to reproduce the observed data satisfactorily. For the five galaxies [A490, A1090, A1190, A1290, A1390], using regularized modeling, the mean model  $\chi^2$  values are [4.59, 1.55, 1.66, 1.58, 1.50] with all galaxies having their maximum individual observable mean  $\chi^2$  values greater than 2.75, and all galaxies having at least three individual observable mean  $\chi^2$  values  $> 1.5$ . Given the high mean  $\chi^2$  values, the model orbit classifications are not displayed here.

In summary, where the kinematical data exhibit bulk rotation of a galaxy, fully aligned rotation is unlikely to be satisfactory: some counter-rotating orbits appear to be necessary to reproduce the observed data. No sense of rotation may be appropriate if no overall rotation or only very minor rotation is observed. This leaves modeling of many galaxies in a position where some overall sense of rotation should be applied to the initial conditions but just how much ought to be determined via experimentation.

#### 5.4. Initial Conditions—An Alternative to 3I

An alternative scheme for initial conditions is the MDJV (match density Jeans velocities) scheme described in Long et al. (2021). The motivation for MDJV is to use the observed data in setting the initial conditions. A set of models of the simulated galaxies using MDJV is constructed and compared with the equivalent 3I baseline models in Section 5.1.

Comparing the MDJV results in Table 7 with the 3I results in Table 3, the mean model  $\chi^2$  values for the MDJV models are all lower than their 3I counterparts, indicating that overfitting has increased. The maximum mean  $\chi^2$  value for MDJV is 0.29 while the maximum for 3I is 0.64. The number of active orbits is increased with the MDJV models. From the results achieved, there is no reason to doubt the validity of using MDJV initial conditions. MDJV does produce orbit classifications that do not agree with those produced by the 3I scheme.

Note that the orbit count circularity classifications for the two schemes are different, with the MDJV models having fewer cold orbits than the 3I models. For the five galaxies [A490, A1090, A1190, A1290, A1390], the cold orbit number percentages are [12.3%, 8.2%, 6.0%, 14.3%, 8.4%] for MDJV and [20.7%, 19.9%, 20.1%, 19.4%, 19.4%] for 3I. It is not obvious that this number difference directly influences the model orbit classification.

#### 5.5. Removal of Orbit Classes

In this section, all orbits of a given orbit circularity class are removed from an existing model and the weighting process rerun. The resulting model is then re-assessed for the impact on reproduction of observables only. Three sets of models are created in this way: the removal of all the cold orbits from the ATLAS<sup>3</sup>D models in Section 5.3, cold orbit removal from the simulated galaxy

baseline models in Section 5.1, and the removal of hot orbits from the simulated galaxy baseline models. To be clear, removing cold orbits means modeling with hot and warm orbits only. Similarly, removing hot orbits means modeling with warm and cold orbits only.

With the removal of cold orbits from the baseline models, for the NNLS regularized models of galaxies [A490, A1090, A1190, A1290, A1390], the mean model  $\chi^2$  values are [0.99, 0.83, 0.47, 1.50, 0.68]. Looking at the individual observable mean  $\chi^2$  values, the observables for galaxy A1290 have not been well reproduced (one observable  $\chi^2$  value  $> 2$ , four  $> 1$ ). Reproduction of observables for the other four galaxies is satisfactory. Removing hot orbits from the baseline models, the mean model  $\chi^2$  values are [1.47, 0.78, 1.06, 0.91, 0.85]. Observables have not been well reproduced for galaxies A490 and A1190; for the other three galaxies, reproduction is satisfactory.

With the removal of cold orbits from the ATLAS<sup>3D</sup> models for galaxies [NGC 1248, NGC 3838, NGC 4452, NGC 4551], the mean model  $\chi^2$  values are [0.34, 0.71, 1.39, 0.21]. Observable reproduction for NGC 4452 is close to unacceptable (all observable mean  $\chi^2$  are  $> 1$ ); there are no concerns for the other three galaxy models.

The models here show a strength of Schwarzschild's method in that it will attempt to weight the orbits it is given to try and reproduce observational data. Whether the orbits are astrophysically realistic, either individually or collectively, is not its concern.

## 5.6. Computer Utilization

The software base for constructing initial conditions and Schwarzschild's modeling is taken from the lead author's implementation of the Syer & Tremaine (1996) made-to-measure stellar dynamical modeling method and the Schwarzschild (1979) orbit based modeling method. This software was first used in Long & Mao (2010), and more recently in Long (2021). Additional analysis code specific to this investigation is written in Python 3 (Van Rossum & Drake 2009). The optimization software used for NNLS is taken from the SCIPY package (Virtanen et al. 2020). The functions used are untailed with default parameter values being taken.

Modeling runs were performed on a 20 core desktop PC. No attempt was made to use graphics processing units (GPUs) to improve performance but their use is not precluded. All software used is Python 3 based with some use of Cython (Behnel et al. 2011) for performance critical code. Multi-processor working, limited to a maximum of 10 processors, was employed. The elapsed times for orbit integration and data collection, and for weight determination (optimization) of the simulated galaxies are shown in Table 8 and in Table 9 for the ATLAS<sup>3D</sup> galaxies. The elapsed times correspond to stages one and two in Section 2.

## 6. Discussion

As set out in the Introduction, Section 1, the science matter is concerned with how to gain knowledge of the internal stellar orbital structure of external galaxies in spite of the inability to obtain 6D phase space data, and to observe stellar orbits. Astronomers have only blended observational data to rely on to try and gain knowledge of structures that cannot be directly validated.

One solution is to try and use an approach such as Schwarzschild's method, and to assume that the orbits that are selected give a good realistic representation of a galaxy, but this is based on assuming that the gravitational potential and the pool of orbits to be selected from have the capacity to represent the galaxy. At best, any such representation can only be indicative or illustrative of the galaxy, and may be significantly inaccurate.

The advent of large language models (LLMs), a type of generative artificial intelligence application, introduces another complication. Astronomers have to be very careful about caveating their results appropriately and not, for example, overstating their results. Astronomy articles form part of the input to LLMs, and new researchers already risk being misinformed by unwisely formulated science articles. This does not mean that Schwarzschild's method is not good at trying to reproduce observational data, but the orbits selected to do so cannot be used for other purposes, unless the selection process has been constructed appropriately. This article offers a partial way forward with orbit classifications being included in the selection process. However, the results here also show that it is clear that Schwarzschild's method cannot be the source of those classifications. Just changing the orbit initial conditions, or using or not using regularization, for example, is sufficient to change the orbit classification for a galaxy.

Using a machine learning application may be an option, if there is some data source which can be used to train the application in time. Perhaps cosmological simulations will become that source. Alternatively, perhaps the reasons why the internal stellar orbital structure is interesting should be re-examined. Where orbital structure is a stepping stone, could it be skipped, for example?

## 7. Conclusions

The objective set out in the Introduction, Section 1, has been met. The question posed was “could an orbit classification for a galaxy being investigated be included in Schwarzschild modeling as a constraint?” As has been demonstrated, the answer is “yes.” Unfortunately, quite where that orbit classification might be obtained from is an open question: it does not appear to be from using Schwarzschild's method given its usual implementation; it might be from cosmological simulations. Having the capability to include classification constraints is, however, another useful technique in the stellar dynamical toolbox.

## Acknowledgments

The author thanks Dandan Xu, Ling Zhu, and Shude Mao for various fruitful discussions as the research recorded here progressed. This work is partly supported by the National Key Basic Research and Development Program of China (No. 2018YFA0404501 to Shude Mao), and by the National Natural Science Foundation of China (NSFC, grant Nos. 11821303, 11761131004 and 11761141012 to Shude Mao).

## ORCID iDs

Richard J. Long <https://orcid.org/0000-0002-8559-0067>

## References

- Behnel, S., Bradshaw, R., Citro, C., et al. 2011, CSE, 13, 31
- Binney, J., & Tremaine, S. 2008, Galactic Dynamics: Second Edition (Princeton, NJ: Princeton Univ. Press)
- Boyd, S., & Vandenberghe, L. 2004, Convex Optimization (Cambridge: Cambridge Univ. Press)
- Cappellari, M. 2002, MNRAS, 333, 400
- Cappellari, M. 2020, MNRAS, 494, 4819
- Cappellari, M., Emsellem, E., Krajnović, D., et al. 2011, MNRAS, 413, 813
- Cappellari, M., Scott, N., Alatalo, K., et al. 2013, MNRAS, 432, 1709
- Chanamé, J., Kleyna, J., & van der Marel, R. 2008, ApJ, 682, 841
- Cretton, N., de Zeeuw, P. T., van der Marel, R. P., & Rix, H.-W. 1999, ApJS, 124, 383
- Emsellem, E., Monnet, G., & Bacon, R. 1994, A&A, 285, 723
- Jin, Y., Zhu, L., Long, R. J., et al. 2020, MNRAS, 491, 1690
- Jin, Y., Zhu, L., Long, R. J., et al. 2019, MNRAS, 486, 4753
- Li, H., Li, R., Mao, S., et al. 2016, MNRAS, 455, 3680
- Long, R. J. 2016, RAA, 16, 189
- Long, R. J., Lu, S.-D., & Xu, D.-D. 2021, RAA, 21, 152
- Long, R. J., & Mao, S. 2010, MNRAS, 405, 301
- Long, R. J., & Mao, S. 2012, MNRAS, 421, 2580
- Long, R. J., & Mao, S. 2018, RAA, 18, 145
- Nelson, D., Springel, V., Pillepich, A., et al. 2019, ComAC, 6, 2
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 2007, Numerical Recipes 3rd Edition: The Art of Scientific Computing (3rd ed.; Cambridge: Cambridge Univ. Press)
- Raissi, M., Perdikaris, P., & Karniadakis, G. E. 2019, JCoPh, 378, 686
- Rix, H.-W., de Zeeuw, P. T., Cretton, N., van der Marel, R. P., & Carollo, C. M. 1997, ApJ, 488, 702
- Rybicki, G. B. 1987, in IAU Symp. Vol. 127, Structure and Dynamics of Elliptical Galaxies, ed. P. T. de Zeeuw (Dordrecht: D. Reidel), 397
- Schwarzschild, M. 1979, ApJ, 232, 236
- Syer, D., & Tremaine, S. 1996, MNRAS, 282, 223
- Tikhonov, A. N. 1963, Sov. Math., 4, 1035

- Valluri, M., Merritt, D., & Emsellem, E. 2004, ApJ, 602, 66  
van den Bosch, R. C. E., van de Ven, G., Verolme, E. K., Cappellari, M., & de Zeeuw, P. T. 2008, MNRAS, 385, 647  
Van Rossum, G., & Drake, F. L. 2009, Python 3 Reference Manual (Scotts Valley, CA: CreateSpace)  
Vasiliev, E. 2013, MNRAS, 434, 3174  
Vasiliev, E., & Valluri, M. 2020, ApJ, 889, 39  
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261  
Xu, D., Zhu, L., Grand, R., et al. 2019, MNRAS, 489, 842  
Zhu, L., van de Ven, G., van den Bosch, R., et al. 2018, NatAs, 2, 233

*Note: Figure translations are in progress. See original paper for figures.*

*Source: ChinaXiv — Machine translation. Verify with original.*