



Application of Convolutional Neural Networks to Identify Stellar Feedback Bubbles in CO Emission

Duo Xu¹, Stella S. R. Offner¹ , Robert Gutermuth² , and Colin Van Oort³

¹ Department of Astronomy, The University of Texas at Austin, Austin, TX 78712, USA; xuduo117@utexas.edu

² Department of Astronomy, University of Massachusetts, Amherst, MA 01003, USA; soffner@astro.as.utexas.edu

³ University of Vermont, Burlington, VT 05405, USA

Received 2019 October 24; revised 2019 December 16; accepted 2019 December 21; published 2020 February 12

Abstract

We adopt the deep learning method called the Convolutional Approach to Shell Identification (CASI) and extend it to 3D (CASI-3D) to identify signatures of stellar feedback in molecular line spectra. We use magnetohydrodynamics simulations modeling the impact of stellar winds in a turbulent molecular cloud to generate synthetic ^{13}CO ($J = 1 - 0$) observations. We train two CASI-3D models: ME1 predicts only the position of feedback, while MF predicts the fraction of the mass coming from feedback in each voxel. We adopt 75% of the synthetic observations as the training set and assess the accuracy of the two models with the remaining data. Both models identify bubbles in simulated data within 5% error. We use bubbles previously visually identified in Taurus in ^{13}CO to validate the models and show that both perform well on the highest confidence bubbles. Models ME1 and MF predict total feedback gas mass of $2894 M_{\odot}$ and $302 M_{\odot}$, respectively. After correcting for missing energy due to the limited velocity range, model ME1 predicts feedback kinetic energies of 4.0×10^{46} erg and 1.5×10^{47} erg with and without subtracting the cloud velocity gradient. Model MF predicts feedback kinetic energies of 9.6×10^{45} erg and 2.8×10^{46} erg with and without subtracting the cloud velocity gradient. Model ME1 predicts bubble locations and properties consistent with previous visual identifications. However, model MF demonstrates that feedback properties computed using visual identifications significantly overestimate feedback impact, due to line-of-sight confusion and contamination from background and foreground gas.

Unified Astronomy Thesaurus concepts: Stellar wind bubbles (1635); Interstellar medium (847); Molecular clouds (1072); Interstellar clouds (834); Astronomy data analysis (1858); Astronomical methods (1043); Convolutional neural networks (1938); Neural networks (1933); Astrostatistics (1882); Interdisciplinary astronomy (804); Stellar feedback (1602); Star formation (1569)

Supporting material: interactive figures

1. Introduction

Stellar winds driven by young stars create distinct features in molecular clouds. The ejected mass compresses and heats the ambient gas, producing shocks (Hollenbach & Tielens 1999). In the case of more massive stars, stellar winds combined with radiation create bubbles containing luminous H II regions. Observational surveys find that signatures of such bubbles are ubiquitous. For example, Churchwell et al. (2006, 2007) visually identified numerous H II regions in the Spitzer Galactic Legacy Infrared Mid-Plane Survey Extraordinaire (GLIMPSE) data. They concluded these bubbles have a significant impact on the dynamics and star formation of molecular clouds, and they found 12% of the shells are associated with young sources, which may have been triggered by shell expansion. Over 50% of the bubbles identified are not spherically symmetric, due to fluctuations in local gas density and/or anisotropic stellar winds and radiation fields. These complications make bubble identification more challenging.

A variety of groups have investigated the impact of stellar feedback bubbles that are due to stellar winds within molecular clouds. Nakamura et al. (2012) found several parsec-scale bubbles expanding and compressing the ambient gas, which they proposed have contributed to the formation of several dense filaments. They suggested that one of these dense filaments is converging with another filament, triggering recent star formation in the cloud. Quillen et al. (2005) mapped the expanding cavities in NGC 1333, a subregion in the Perseus

molecular cloud. They found the kinetic energy of these cavities is sufficient to power the turbulence in this region. Arce et al. (2011) identified stellar feedback bubbles in a full map of the Perseus molecular cloud and drew similar conclusions about the energy budget. However, Li et al. (2015) found that the energy injected from bubbles in the Taurus molecular cloud is only 29% of the turbulent energy of Taurus.

A variety of theoretical works have also investigated the impact and signatures of stellar feedback. Boyden et al. (2016) studied statistical signatures of stellar winds in synthetic observations. They found that the covariance matrices of the velocity channels are sensitive to the existence of stellar feedback. Offner & Liu (2018) found that the slope of the velocity power spectrum becomes steeper in simulations with feedback compared to those without feedback. However, stellar feedback might also indirectly add energy to the ambient gas where there is no injected feedback mass (Offner & Liu 2018), making quantitative statistical study of feedback more challenging in observational data.

Historically, bubbles, such as those found in the above studies, have been identified “by eye.” However, given the exponentially increasing amount of data, visual identification is not scalable; that is, it is almost impossible for humans to look through all the data by eye (Molinari et al. 2010; Peek et al. 2011). Moreover, to study the dynamics of bubbles, it is necessary to switch from two-dimensional images (Beaumont et al. 2014) to three-dimensional

data cubes (Arce et al. 2011; Li et al. 2015) because those contain information about the gas motion. An extra dimension makes it much more time intensive to identify the bubble features. Moreover, it is impossible for humans to consistently identify or classify without bias. However, systematic and repeatable identification is possible with the aid of machine-learning approaches (Beaumont et al. 2011, 2014; Van Oort et al. 2019).

Several machine-learning algorithms have been applied to identify stellar feedback features (Beaumont et al. 2011, 2014). Beaumont et al. (2011) applied support vector machines (SVM) to distinguish a supernova remnant from the ambient gas in $^{12}\text{CO } J = 3 - 2$ emission. Beaumont et al. (2014) developed the Brut algorithm, based on random forests, to identify bubbles in dust emission. To train Brut, they adopted bubble identification results from over 35,000 participants in the Milky Way Project, a citizen science project based on GLIMPSE data. Xu & Offner (2017) expanded on this work by supplementing the Brut training set with synthetic observations of bubbles in simulated clouds. After retraining on the enhanced training set, Brut more efficiently identified ultracompact and compact H II regions generated by B-type stars. Leveraging both observational data (e.g., Simpson et al. 2012; Jayasinghe et al. 2019) and synthetic observations can significantly enhance the performance of machine-learning algorithms. However, Brut requires the bubble to be centered in the image for it to be accurately identified. This makes it computationally expensive to identify bubbles in a large sky survey map because the data must be cropped into small chunks centered at different positions with different sizes to ensure the target bubbles are centered in at least one image.

Due to the evolution of high-performance computing and the power of graphics processing units (GPUs), deep learning is gaining popularity thanks to its general applicability and high accuracy. Recently developed deep learning methods are more powerful in image recognition than earlier methods, like Brut. Ntampaka et al. (2019) developed a deep machine-learning tool based on convolutional neural networks (CNNs) to estimate the mass of galaxy clusters in X-ray emission. The CNN is not sensitive to the position of galaxy clusters, making it straightforward to apply it to large sky survey maps. Van Oort et al. (2019) developed an “encoder-decoder” Convolutional Approach to Shell Identification (CASI) to identify stellar wind bubbles in density slices and 2D CO emission. Once trained, CASI can identify structures in minutes and achieves a 98% pixel-level accuracy. However, one caveat of these CNN models is that they are limited to 2D integrated intensity maps or individual velocity channels. These algorithms do not take radial velocity information into consideration, which may lead to a high false detection rate when applied to 3D data. In other words, this technique may identify a clear ring structure as a bubble even though this structure is caused by a turbulent pattern without any evidence of expansion in the spectra. Alternatively, it may miss structures with coherent expansion across a range of channels but which do not have a bubble morphology in most channels.

In this paper, we adopt the deep learning method CASI (Van Oort et al. 2019) and extend it to 3D (CASI-3D) in order to exploit the full 3D CO spectral information to identify bubbles. We develop two deep machine-learning tasks. Task I predicts the position of feedback. Task II predicts the fraction of the mass in the voxels that is coming from feedback. We describe our deep machine-learning algorithm architecture and how we

generate the training set from synthetic observations in Section 2. In Section 3, we present the performance of the CNN model in identifying bubbles in both synthetic data and observational data. We summarize our results and conclusions in Section 4.

2. Method

2.1. CASI-3D Architecture

In this section, we give a brief overview of the CASI architecture and describe our implementation in 3D. CASI-3D is available on GitLab.⁴

Van Oort et al. (2019) developed the CASI architecture with residual networks (He et al. 2016) and “U-net” (Ronneberger et al. 2015). The residual networks exponentially increase the complexity of the networks to prevent overfitting (He et al. 2016; Veit et al. 2016). “U-net” adds cross connections between different layers, which enhance the performance in constructing the output image (Ronneberger et al. 2015). CASI utilizes a widely used optimization method, stochastic gradient descent (SGD) with momentum, in the training. Momentum indicates that the SGD takes the past time step into consideration when conducting the optimization. It helps accelerate SGD during training. CASI is trained on simulated molecular cloud density 2D slices or CO integrated intensity maps. It learns to predict the “tracer field,” which is an ORION field (Li et al. 2012) that tracks the fraction of gas in each cell that is launched in the stellar winds. A more detailed description of CASI can be found in Van Oort et al. (2019).

We modify the CASI architecture and replace the 2D convolutions with 3D convolutions, as shown in Figure 1. The “encoder” part extracts the features from the input data, and then the “decoder” part translates these features into another image, such as the tracer field. In Van Oort et al. (2019), the input data is a 256×256 array. The image goes through four down-sampling (max pooling) layers. Given the extra dimension in our model, we cannot maintain the same spatial resolution in our 3D model, due to the limited memory on the GPU. We re-form the input data to an array shape of $64 \times 64 \times 32$ in position–position–velocity (PPV). The data cube undergoes three down-sampling layers. We apply $7 \times 7 \times 7$ filters in the first layer to represent the large-scale structure in the training set. We then apply $3 \times 3 \times 3$ filters in the following layers, which capture the details of bubble morphology and velocity information. We begin with 32 filters and double them after each down-sampling layer. We adopt the average pooling method to down-sample the data. See Appendix A.1 for more discussion of the down-sampling methods.

Owing to the limitation of the GPU memory, we set the batch size to be eight. The batch size indicates the number of samples that propagate through the network at one time. We explore different numbers of epochs during the training. An epoch is one iteration of the prediction, error estimation, and weight update cycle over all the training data. After some number of epochs, the model error converges and does not decrease much more. The learning rate is a hyperparameter that controls the magnitude of filter weight updates with respect to the loss gradient. The learning rate affects how quickly our model converges to an acceptable minimum squared error

⁴ <https://gitlab.com/casi-project/casi-3d>




Figure 1. Architecture of the U-net CNN model with residual functions.

value. If the learning rate is too high, the model diverges and cannot reach the minimum squared error value. If the learning rates are too low, it will take a long time for the model to converge as it makes very tiny updates to the weights in the network.

We test two strategies to set the learning rate during training: a fixed learning rate during the whole training process or an adaptive learning rate that changes with the epoch. The fixed learning rate is set to 0.01. The adaptive learning rate is initially set to 0.02 and halved every 19 epochs. Considering the computing time limits during the training, we set the maximum epoch number to be 277. Meanwhile, we set an early-stopping criterion for models ME1, ME2, ME5, ME6, and ME7. The early-stopping criterion provides guidance on how many epochs can be run before the model begins to overfit. The model stops the training if the validation loss decreases less than 0.0001 after 100 epochs. In practice, most models (ME2, ME5, ME6, ME7) reach the early-stopping criterion before reaching the epoch limit.

2.2. Training Task

We develop two training tasks to identify stellar feedback bubbles. Task I involves training to predict the position of feedback, which reproduces the morphology of bubbles. Task II involves training to predict the fraction of the mass that comes from stellar feedback, which gives a better mass estimation of the bubbles.

2.2.1. Training Task I: Predicting the Position of Feedback

In Task I, we aim to predict the position of feedback on a pixel-by-pixel basis. We use the mean squared error (MSE) as the loss function in the training. The loss function is a metric to quantify the performance of the model predictions. The mean squared error is defined as

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (Y_{\text{pred}} - Y_{\text{tracer}})^2, \quad (1)$$

where Y_{pred} represents the prediction from the model, Y_{tracer} represents the “ground truth” as described by the tracer field, and n indicates the number of samples. In this model, Y_{tracer} is the emission of ^{13}CO at the specific positions where there is

feedback gas. The ^{13}CO intensity is proportional to the mass of the gas in the optically thin regime.

Observational data includes noise, so we must adopt a prescription to define feedback that accounts for noise limits. We adopt MSE in part because it is less affected by the noise in the data cube than other loss functions. We set a threshold based on the input data noise level, 0.2 K, to binarize the prediction map; that is, a pixel is assigned a 1 if the predicted emission is above 0.2 K, and a pixel is assigned a 0 if the predicted emission is below 0.2 K. The binarized prediction map indicates the location of the bubbles in PPV space. The final loss indicates the error of the prediction, which is related to the mass estimation uncertainty.

To explore the performance of different models with different hyperparameters, we train models with different learning rates, different epochs, different optimization methods, and different loss functions. We easily rule out other loss functions and optimization methods based on the performance on the training set, and we adopt MSE as the loss function and SGD as the optimization method. We list different models with different learning rates and with different epochs in Table 1. We additionally explore the completeness of the training set on the performance. More discussion can be found in Section 2.3.4.

2.2.2. Training Task II: Predictions for the Fraction of Feedback Mass

Task I basically classifies pixels individually as belonging to feedback or not belonging to feedback. However, it does not take into account that many pixels contain some feedback and some nonfeedback gas. To address this, we train another model to predict the fraction of the mass that comes from stellar feedback. We adopt the same learning rate, epoch, and optimization method as the best model in Section 2.5. We define the new training data to be PPV cubes containing the fraction of mass, rather than the emission, that comes from the feedback. We describe how these cubes are constructed in more detail in Section 2.3.2.

We test the MSE, intersection over union (IoU), and a combination of MSE and IoU as the loss function and compare their performance in Appendix A.2. The IoU is a metric to evaluate the overlap fraction between the prediction and the

Table 1
Training Model Parameter

Model	Task	Training Set			Learning Rate	Epoch
		Fiducial Resolution (5 pc $\times$ 5 pc)	Negative Set	High Resolution (2.5 pc $\times$ 2.5 pc)		
ME1	I	✓	✓	✓	adaptive	277
ME2	I	✓	✓	×	adaptive	223
ME3	I	✓	✓	✓	fixed	60
ME4	I	✓	✓	✓	adaptive	60
ME5	I	✓	×	✓	adaptive	189
ME6	I	✓	×	✓	adaptive	260
ME7	I	×	✓	✓	adaptive	260
MF ^a	II	✓	✓	✓	adaptive	277

Note.

^a The training set and the hyperparameters of model MF are the same as those of ME1. The only difference is the training set. Model MF adopts the fraction of the feedback mass as the target. The other models adopt the intensity of the feedback emission in the training.

tracer field and is defined as

$$\text{IoU} = \frac{Y_{\text{pred}} \cap Y_{\text{tracer}}}{Y_{\text{pred}} \cup Y_{\text{tracer}}}. \quad (2)$$

The combination of MSE and IoU is defined as

$$L_{\text{new}} = \omega \times \text{MSE} + \text{IoU}, \quad (3)$$

where ω is the weight of the MSE in the new loss function. Here we set $\omega = 100$. Since the fraction is between 0 and 1, the MSE is not sensitive to small values, while the IoU is strongly sensitive to small values. We find that the combination loss function performs the best in predicting the fraction of the mass that comes from stellar feedback.

2.3. Training Sets

2.3.1. Synthetic ^{13}CO Observations

To train CASI-3D, we adopt magnetohydrodynamic (MHD) simulations that model sources launching stellar winds in a piece of a turbulent molecular cloud (Offner & Arce 2015). The simulation box is $5 \times 5 \times 5 \text{ pc}^3$, with a total mass of $3762 M_{\odot}$ and mean density of $\sim 500 \text{ cm}^{-3}$. Offner & Arce (2015) conduct different simulation runs with different mass-loss rates, different magnetic fields, different turbulent patterns, and different evolutionary stages to study the impact of stellar winds on the ambient gas. More details about the simulations can be found in Offner & Arce (2015).

We apply the publicly available radiation transfer code RADMC-3D (Dullemond et al. 2012) to model the ^{13}CO ($J = 1 - 0$) line emission of the simulation gas. We use ^{13}CO emission rather than ^{12}CO because ^{12}CO is optically thick at the average simulation column density. To construct the synthetic observations, we adopt the simulation density, temperature, and velocity distribution for the RADMC-3D inputs. In the radiative transfer, H_2 is assumed to be the only collisional partner with ^{13}CO . In general, we assume the ^{13}CO abundance is constant where $[^{13}\text{CO}/\text{H}_2] = 1.5 \times 10^{-6}$. However, when gas conditions are likely to result in full dissociation of ^{13}CO ($T > 1000 \text{ K}$ or $n(\text{H}_2) < 50 \text{ cm}^{-3}$), we set the abundance to zero. In addition, we also set the ^{13}CO abundance to zero in conditions where it would freeze out onto dust grains ($n(\text{H}_2) < 10^4 \text{ cm}^{-3}$) or where it would be dissociated by strong shocks ($|v| > 10 \text{ km s}^{-1}$).

We increase the training set by also considering thin clouds. Qian et al. (2015) study the thickness of molecular clouds from the core velocity dispersion (CVD) and find that the line-of-sight thicknesses of Taurus, Perseus, and Ophiuchus molecular clouds are $\leq \frac{1}{8}$ of their length. To recreate the distinct circular structures of the observed stellar feedback, we crop the data cube into thinner slices, including widths of 2 and 0.9 pc. We show the emission from ^{13}CO with different cloud thicknesses in Appendix A.1 and discuss why we adopt ^{13}CO in favor of ^{12}CO in our analysis.

2.3.2. Training Target: Tracer Field for Task I

To construct the training set for Task I, we first define the position of the bubbles using the tracer field that indicates the fraction of gas coming from stellar feedback at each position. A given voxel is assigned to be part of a bubble structure where more than 2% of the mass comes from stellar feedback. Conversely, a voxel is assigned to be pristine gas where less than 2% of the mass comes from stellar feedback. To better capture the morphology of the bubbles, we define the position of a bubble by the gas temperature, where $T \geq 12 \text{ K}$. We discuss different definitions of bubbles in Appendix A.2, including a criterion using the gas velocity.

We mask all the positions of pristine gas. We set the ^{13}CO abundance to be 0 in the masked region and compute the radiative transfer to obtain the ^{13}CO emission that is only coming from the stellar feedback, which we refer to as the ^{13}CO feedback map. Figure 2 shows an example of synthetic ^{13}CO observations and the ^{13}CO feedback map. In PPV space, the feedback map is the ^{13}CO emission that comes from stellar feedback, which allows us to distinguish between the feedback and diffuse emission from the host molecular cloud. The wind tracer is the positive signal that CASI-3D learns to pick out from the messy ^{13}CO emission of the molecular cloud.

In PPV space, the ^{13}CO feedback map emission is only a fraction of the total emission in each voxel, due to the foreground and background emission along the line of sight. As a result, we built the target in Task I by filling the voxels where there is feedback with the corresponding ^{13}CO emission in PPV space. This closely emulates what astronomers do in observational data (Arce et al. 2011; Li et al. 2015). The model in Task I adopts the raw ^{13}CO data cube and returns a data cube of the same shape that only has the ^{13}CO intensity in the feedback regions.




Figure 2. Upper row: ^{13}CO integrated intensity including emission from the full cloud. Bottom row: integrated intensity of the full ^{13}CO (upper row) masked by a ^{13}CO synthetic observation of the tracer field to obtain the pixel locations of the feedback in PPV space (see Section 2.3.2). First column: synthetic observations corresponding to model W2_T2_t0 in Offner & Arce (2015). Second column: synthetic observations corresponding to model W2_T2_t1. Third column: synthetic observations corresponding to model W1_T2_t0. Fourth column: synthetic observations corresponding to model W2_T3_t0.

In the final step, we compute the bubble properties from the output prediction. The prediction is the reconstructed ^{13}CO emission in voxels associated with feedback in PPV space. We apply a threshold based on the noise level of the input data to mask out noise-induced false bubble identifications. We then sum over all the voxels to calculate the mass, momentum, and energy of the identified bubbles.

2.3.3. Training Target: Mass Fraction for Task II

To attempt to improve mass, momentum, and energy estimation in the feedback identification, we adopt the fraction of feedback mass as the target in Task II. To build the target, we first convert the raw density from position–position–position space to PPV space, just like the ^{13}CO data cube. Next, we convert to PPV space the density cubes that only include the feedback gas. We take the ratio of the two converted cubes to get the fraction of feedback mass in PPV space. The fraction ranges from 0 to 1, and the fraction value is not necessarily proportional to the actual ^{13}CO intensity. If the ^{13}CO emission is optically thin, its column density is proportional to the emission intensity. Knowing the fraction of the feedback in each position allows us to calculate the actual feedback mass.

Note that the emission predicted by Task I does not exclusively come from the feedback gas. Pristine gas that is not associated with feedback also contributes to the emission. Thus, Task I overestimates the mass coming from feedback. Consequently, Task II is a more advanced approach to estimating the physical properties of the feedback bubbles.

2.3.4. Data Augmentation

We adopt simulations with different mass-loss rates, different magnetic fields, different turbulent patterns, and different evolutionary stages to create synthetic observations. The simulations have a physical scale of 5 pc on a side. To enhance the diversity of the training set, we conduct radiative

transfer from three different angular views and rotate the images every 15° from 0° to 360° . Figure 2 shows an example of the integrated intensity and the tracer field of ^{13}CO with different model outputs. We also construct a “zoomed-in” synthetic observation on bubbles with an image size of $2.5 \text{ pc} \times 2.5 \text{ pc}$. Different scales of bubbles in the training set reinforce the ability of the model to detect bubbles on different scales.

To help CASI distinguish feedback bubbles from shell-like structures produced from supersonic turbulence in the molecular cloud, we also conduct synthetic observations on purely turbulent simulations including noise, which do not contain feedback sources. We adopt the nonfeedback cloud emission data as a negative training set. This negative training set is essential because it trains the algorithm to ignore large (e.g., 0.5–2 pc) arc-like and shell-like features that ambient turbulent motions may generate despite no recent stellar feedback. In the Taurus survey, for instance, there are areas that are considerably larger than our adopted postage stamp view-port size that lack any feedback sources while curving, filamentary structures are visually apparent. The model must perform well in those areas or else may provide false detections. Table 1 lists the properties of the training sets adopted by seven different models.

To make the synthetic cubes closer to the real observational data in Section 2.4, we convolve them with a telescope beam of $50''$ and add noise. We assume the synthetic images are at a distance of either 140 or 250 pc and are observed by Five College Radio Astronomy Observatory (FCRAO; Ridge et al. 2006). Figure 3 shows a bubble before and after we convolve the image with the beam and add noise. The noise level, 0.125 K, is the same as the rms noise in the Taurus ^{13}CO observational data. Moreover, we randomly shift the central velocity of the cubes between -1 and 1 km s^{-1} to increase the diversity of the training set.

In total, we generate 7821 synthetic data cubes: 3910 have a field of view (FoV) of $5 \text{ pc} \times 5 \text{ pc}$, 3648 have an FoV of




Figure 3. Integrated intensity of ^{13}CO ($J = 1 - 0$). Left: ^{13}CO integrated intensity of a simulated bubble without noise. Middle: ^{13}CO integrated intensity of a simulated bubble convolved with a beam of size $50''$ and with 0.125 K noise. Right: ^{13}CO integrated intensity of an observed bubble, TMB 8, identified by Li et al. (2015).

$2.5 \text{ pc} \times 2.5 \text{ pc}$, and 509 contain no feedback sources. We adopt 4693 of the data cubes as a training set, 1564 data cubes as a test set, and 1564 data cubes as a validation set. The validation set allows us to estimate how well the model has been trained. The test set assesses the accuracy of the final model.

2.4. Taurus Data

The Taurus ^{13}CO $J = 1 - 0$ map was observed between 2003 and 2005 using the 13.7 m FCRAO Telescope (Narayanan et al. 2008). The map covers an area of $\sim 98 \text{ deg}^2$ with a beam size of $50''$. The data have a mean rms antenna temperature of 0.125 K. We combine this with the Class III young stellar object (YSO) catalog of Kraus et al. (2017) as a reference to determine the potential driving sources of the bubbles. Kraus et al. (2017) reexamined 396 candidate members from previous surveys in the literature covering $3^{\text{h}}50^{\text{m}} < \alpha < 5^{\text{h}}40^{\text{m}}$ and $14^\circ < \delta < 34^\circ$. They concluded 218 YSOs are confirmed or likely Taurus members, but 160 candidates are confirmed or likely interlopers, and the remaining 18 objects are uncertain.

Li et al. (2015) visually identified 37 bubbles in the Taurus molecular cloud from the ^{13}CO emission, and we adopt these as an observational test sample for our models. We divide these 37 bubbles into three categories based on their morphology and likely driving source. The three ranked categories of bubbles are as follows:

- A. An A bubble contains at least one YSO inside the bubble and has a clear circle/arc morphology.
- B. A B bubble contains no YSOs inside but contains at least one YSO on the bubble rim or near the bubble boundary.
- C. A C bubble contains no known YSOs in or around the bubbles.

Among the 37 bubbles, 7 are Rank A, 16 are Rank B, and 14 are Rank C.

To make the observational data suitable for the algorithm, we first down-sample the Taurus ^{13}CO data cube by a factor of 3 so that it has a resolution ($\sim 1'$) similar to that of the training set. We shift the down-sampled cube's mean velocity to 0, and we then crop the velocity range from -4 to 4 km s^{-1} . We crop the cube to $2.7 \text{ pc} \times 2.7 \text{ pc}$ after centering on each bubble identified by Li et al. (2015). This procedure generates a stack of data cubes with a shape of $64 \times 64 \times 32$. See Appendix D for more detail.




Figure 4. Training and validation errors of model ME1 during training.

However, investigating only the 37 previously identified bubbles is limiting. CASI-3D does not require the bubbles to be centered in the image, and the algorithm is able to rapidly search the entire Taurus map if it is divided into smaller data cubes. Furthermore, comparing the CASI-3D identification of the previously defined, cropped bubbles to those identified in the full map in an unbiased search allows us to verify that the algorithm is translation invariant and insensitive to the position of the bubble.

CASI-3D requires input data that has the same dimensions as the training data. When applying the CNN model to the full map, we decompose the Taurus data into a series of $64 \times 64 \times 32$ cubes. Each cube is offset by five pixels ($\sim 5'$), resulting in 92% overlap with adjacent steps. We begin cropping from the northeast corner of the full map. Then we move to the next position along the R.A. direction with a step size of five pixels. When we finish the sampling procedure along the R.A. direction at fixed decl., we move five pixels in decl., and then repeat the process again.

2.5. Model Selection

2.5.1. Validation

After training, we find all models in Task I converge to an MSE below 0.1, and a combination of MSE and IoU in Task II (model MF) converges to unity. Figure 4 shows the training




Figure 5. Integrated intensity of the ^{13}CO (top left), the integrated intensity of the full ^{13}CO masked by a ^{13}CO synthetic observation of the tracer field to obtain the pixel locations of the feedback in PPV space (second from top left), and the prediction of eight CNN models for the bubble emission.

and validation errors of model ME1. After 277 epochs, this model converges to an MSE of 0.06. The number of epochs used in the training, 277, is set by the maximum job run-time permitted on our computing resources. We show the performance of seven CNN models on a test set of synthetic observations in Figure 5. All models in Task I and II clearly capture the shell features produced by stellar feedback.

2.5.2. Mean Opinion Score: Visually Assessing the Model Performance

To visually assess the performance of the CNN models on observational data, we apply eight CNN models as listed in Table 1 to the Taurus ^{13}CO bubble data. All test bubbles have a clear circular or arc-like structure across a range of velocity channels, which provides an appropriate test sample. However, we do not quantitatively know the true bubble boundaries, so we use visual identification and assessment to evaluate the performance of the models. We introduce the mean opinion score (MOS) to visually quantify the performance of the Task I models on the Rank A and B bubbles in Taurus. The MOS is expressed as an integer ranging from 1 to 5, where 1 is poor performance and 5 is excellent performance. We create a rubric outlining the characteristics of each score to ensure that visual ranking between assessors is as uniform as possible. The rubric is as follows:

1. Excellent performance. The prediction covers the full rim structure of the bubble, with less than 10% extra emission, that is, a minimal amount of false-positive pixels.
2. Good performance. The prediction covers 75% of the rim structure of the bubble, with less than 25% extra emission.
3. Average performance. The prediction covers half the rim structure of the bubble, with less than 50% extra prediction.
4. Fair performance. The prediction covers one-third of the rim structure of the bubble, with less than 70% extra prediction.

5. Poor performance. The prediction covers less than one-third of the rim structure of the bubble, with more than 70% extra prediction.

Figures 6 and 7 show the integrated intensity of example Rank A and B bubbles and compare the predictions from the CNN models.

We conduct a blind rating of the performance of the seven models in Task I, where each coauthor assessed the quality of each bubble prediction. To judge the prediction, the coauthors looked at the integrated intensity map and the channel-by-channel prediction of each model for each bubble. Figure 8 shows the MOS of each model as rated by the four coauthors. Five of the seven models have an overall MOS that is above average. We easily rule out models ME6 and ME7, which have lower MOSs. ME1 exhibits decent performance on the Taurus bubbles, and it is trained using the most complete training set that includes both negative examples (no bubbles) and higher resolution bubbles. We adopt ME1 as the fiducial model in the following analysis.

For Task II, we maintain the same training set and hyperparameters as those adopted for ME1 and simply replace the training target data with the fraction of mass coming from feedback in the training set, as discussed in Section 2.1.

3. Results

3.1. Assessing Model Accuracy Using Synthetic Observations

In this section, we use the synthetic images to assess how accurately physical properties can be determined from the identified bubbles. We apply all of the models to the synthetic observations of the bubbles in the test set, as shown in Figure 5. We mask the prediction cubes at 0.2 K, which is consistent with the noise level of the input cubes. We then calculate the mass of the bubbles by assuming that the ^{13}CO emission line is optically thin and has an excitation temperature of 25 K (Narayanan et al. 2012; Li et al. 2015). We examine the uncertainty of the bubble mass estimation in terms of the choice of excitation temperatures in Appendix C. We take 1.5×10^{-6} as the abundance ratio between ^{13}CO and H_2 (Narayanan et al. 2012; Li et al. 2015).




Figure 6. Results of the eight models applied to Rank A bubble TMB12. First row, first column panel: integrated intensity of ^{13}CO overlaid with a yellow ring indicating the position and thickness of the bubble. Star symbols show the location of the Class III YSOs from Kraus et al. (2017). The remaining panels show the predicted intensity integrated along the velocity axis for the eight CNN models.

Finally, we compute the total mass by summing over the bubble volume.

Figure 9 shows the mass estimated from the two models, ME1 and MF. We also plot the true feedback mass, which we estimate directly by adding the mass contained in all cells with $T \geq 12$ K and a tracer fraction $\geq 2\%$ (see Section 2.3.2). We find ME1 overestimates the bubble mass by a factor of 3 or more, while MF correctly predicts the bubble mass within 4% error. The low-mass bubbles are overestimated by a factor of 10 by model ME1, while the high-mass bubbles are overestimated by a factor of 3 by model ME1. Since the low-mass bubbles are usually small, they do not expand enough to break out of the cloud such that more gas along the line of sight contributes to the ^{13}CO emission. Their velocities are also small, yielding more surrounding gas in the velocity channels where the feedback is. On the other hand, high-mass bubbles are usually large with more gas coming from the driving YSOs and have larger expanding velocities. The gas along the line of sight of the high-mass bubbles occupies a smaller fraction (but still a large portion) in each velocity channel compared to that of low-mass bubbles.

We compare the 1D line-of-sight momentum between the model prediction and the true simulation feedback in Figure 10. We define the 1D momentum as the sum of the gas mass in each channel multiplied by the channel velocity, where we have shifted the mean cloud velocity to zero. Model ME1 overestimates the 1D momentum by a factor of 2.8. In contrast, model MF is able to correctly predict the 1D momentum within 10% error.

Under the assumption of isotropic expansion, the 3D momentum would be expected to be a factor of $\sqrt{3}$ larger than the 1D momentum, while the 3D kinetic energy would be a factor of 3 larger than the 1D kinetic energy. Figures 11 and 12 show the 1D momentum and energy, respectively, predicted by the two models compared to the respective 3D quantities calculated from the simulation. Again, we find that the momentum and energy predicted by model MF are comparable to the true simulation values.

One caveat here is that we limit the velocity range of the synthetic observations in the training set to match the Taurus observation. To assess how much mass, momentum, and energy are missed by applying this cutoff, we calculate the total mass, momentum, and energy associated with velocities that exceed the CO spectrum velocity range. We find that 9% of the mass is in gas with $|v| > 4 \text{ km s}^{-1}$. Meanwhile, 20% of the momentum and 44% of the energy are missed because of the limited CO velocity range. In observations, the amount of missing mass, momentum, and energy will depend both on the observation spectral range and the source masses, which are often not well constrained. In the following sections, we correct the totals for the missing mass, momentum, and energy.

The different accuracies achieved by the ME1 and MF models can be understood as follows. Model ME1 is trained using the tracer intensity and thus can only predict the feedback position. We find that the training set for ME1, namely the full ^{13}CO emission masked by the position of the tracer field in PPV space, is not a good indicator of the fraction of feedback mass in each voxel, because the emission at the predicted position does not exclusively come from feedback gas. Instead, gas along the line of sight that is not associated with feedback contributes to the emission and can dominate the total. This matches the current state of the art in human identification, because visual identification of feedback cannot disentangle the feedback from the nonfeedback gas within a given velocity range and voxel. This is why our model ME1 predicts total mass, momentum, and energy similar to that estimated by Li et al. (2015), as we will show in Section 3.2. However, model MF adopts the fraction of mass coming from feedback as the training target, which allows the model not only to predict the position of feedback but also to predict the fraction of the mass coming from feedback in each voxel. This allows a significantly more accurate determination of the mass, momentum, and energy.

3.2. Physical Properties of the Individual Taurus Bubbles

Next, we estimate the masses of the bubbles in Taurus identified by models ME1 and MF and compare them with the




Figure 7. Same as Figure 6 but depicting the Rank B bubble TMB37.




Figure 8. Overall mean opinion score (MOS for each model, averaged over all 23 Rank A and Rank B bubbles and the visual rankings by the four human judges).

previous observational estimates. For the purpose of comparison with the prior visual identifications, we analyze postage stamps centered on each of the 37 Rank A, B, and C bubbles. We calculate the bubble mass and momentum for each model as described in Section 3.1. The observational approach to calculating the observed bubble mass, momentum, and energy is as follows. Li et al. (2015) adopt an annulus as a mask for each bubble rim region, where the inner and outer radii of the annuli are determined by visual inspection, and then add up the emission of ^{13}CO in the masked region to calculate the mass. The bubble velocity extent along the line of sight is also determined by eye.

Figure 13 compares the bubble mass calculated from the two CNN models and that from the observational approach for all the bubbles identified by Li et al. (2015). The mass estimated from ME1 shows an approximately linear trend compared with that from the observational approach within a factor of 2. Figure 13 also compares the mass calculated from MF with the mass estimated from the observational approach. We find the




Figure 9. Relation between the CNN-predicted bubble mass and the true feedback mass for different bubbles. The filled symbols indicate the mass calculated from model MF. The open symbols represent the mass calculated from model ME1. The black dashed line indicates where the CNN correctly predicts the true mass as determined by the tracer field and gas temperature. The blue dashed line has a slope of 3, and the red dashed line has a slope of 10. The physical parameters of the simulations with different labels are listed in Offner & Arce (2015).

observational approach overestimates the bubble mass by an order of magnitude.

Figures 14 and 15 compare the momentum and energy calculated from the CNN models with those from the observational approach. The momentum and energy estimated from ME1 both show approximately linear trends compared to those from the observational approach and are within a factor of 2. However, when considering the fraction of mass coming from feedback, both model ME1 and the observational approach overestimate the momentum and energy by an order of magnitude.

The differences between the observational approach and model ME1 are not too surprising because the observational approach to calculating the bubble mass, momentum, and energy is fairly simple. For example, Li et al. (2015) may overestimate or underestimate the mass depending on the ring




Figure 10. Relation between the CNN-predicted bubble momentum and the true feedback 1D momentum from different bubbles. The filled symbols indicate the momentum calculation from model MF. The open symbols represent the momentum calculation from model ME1. The dashed line indicates where the CNN correctly predicts the true momentum.




Figure 11. Relation between the CNN-predicted bubble momentum and the true feedback 3D momentum from different bubbles. The filled symbols indicate the momentum calculation from model ME1. The open symbols represent the momentum calculation from model ME1. The black dashed line has a slope of 1, and the blue dashed line has a slope of $1/\sqrt{3}$, which indicates that velocity symmetry is a reasonable assumption in estimating the true 3D momentum.

mask they draw. In addition, bubbles are not always a closed, symmetric circle. They are likely to be discontinuous on the rim, as depicted in Figure 6. As we showed in Section 3.1, the higher values obtained by model ME1 and the observations compared to model MF are also not surprising because the former approaches include excess material along the line of sight that is not part of the feedback. However, the CASI-3D models may also overestimate the total bubble properties because there may be more than one bubble identified in each postage stamp. To address this, we set the postage stamp size to minimize this effect.

Several different effects may cause errors in estimating the mass, momentum, and energy from the CO emission. We find that the choice of excitation temperature could cause a factor of 2 error in mass estimation, but it cannot account for a factor of




Figure 12. Relation between the CNN-predicted bubble energy and the true feedback energy from different bubbles. The filled symbols indicate the energy calculation from MF. The open symbols represent the mass calculation from ME1. The black dashed line has a slope of 1, and the blue dashed line has a slope of one-third.




Figure 13. Bubble mass estimated from the CASI-3D model predictions and the observational mass estimate from Li et al. (2015). The gray dashed line indicates the trend for equal mass, while the purple dashed line is scaled down by 10. The blue symbols indicate the mass calculated from model ME1. The red symbols represent the mass calculated from model MF.

10 (see Appendix C for more detail). Likewise, the assumption of LTE has a small effect on the mass estimation. We conclude that the line-of-sight gas contamination is the main uncertainty in mass estimation. As we discussed in Section 3.1, low-mass bubbles are overestimated by a larger factor (a factor of 10) than high-mass bubbles (a factor of 3), due to the line-of-sight gas contamination. For low-mass bubbles, the line-of-sight contamination is the dominant factor overestimating the mass, but for high-mass bubbles, the uncertainty that comes from line-of-sight contamination is similar to the uncertainty that comes from assuming a fixed excitation temperature.

It is also worth considering the estimates from a physical perspective. The mass associated with feedback based on our understanding of the launching velocities of feedback is usually comparable to the young star's mass. It is physically impossible for a few $M_\odot$ young star to drive a 50–60 $M_\odot$ bubble. Arce et al. (2011) pointed out that these high-mass bubbles could be produced




Figure 14. Bubble momentum estimated from the CASI-3D model predictions and the observational momentum estimate from Li et al. (2015). The gray dashed line indicates the trend for equal momentum, while the purple dashed line is scaled down by 10. The blue symbols indicate the momentum calculated from model ME1. The red symbols represent the momentum calculated from model MF.




Figure 15. Bubble energy estimated from the CASI-3D model predictions and the observational energy estimate from Li et al. (2015). The gray dashed line indicates the trend for equal energy, while the purple dashed line is scaled down by 10. The blue symbols indicate the energy calculated from model ME1. The red symbols represent the energy calculated from model MF.

if $v_{\text{wind}} = 200 \text{ km s}^{-1}$ with $\dot{m}_{\text{wind}} = 10^{-7}\text{--}10^{-6} M_{\odot} \text{ yr}^{-1}$. However, these mass-loss rates are orders of magnitude higher than those that can be explained by stellar winds, outflows (considering outflows are collimated; Bally 2016), or ionization or radiation pressure from the stars observed in Taurus (Smith 2014). The mass directly launched by young stars in both theoretical work and observations is small, $\sim 10^{-9} M_{\odot} \text{ yr}^{-1}$ (e.g., Shu et al. 1994; Hartigan et al. 1995). Numerical simulations suggest the entrained gas can contribute three times more mass than the direct mass loss from young stars (Offner & Chaban 2017). In observations, the mass associated with feedback is included in the estimate of the entrained gas. Model ME1 does the same thing: it predicts the gas associated with feedback (including the entrained gas) but cannot disentangle the line-of-sight contamination. Model MF goes one step further to predict the fraction of gas mass associated with feedback (including the entrained gas). Although we include the entrained gas in the bubble mass estimation, the result is significantly less than $10\text{--}100 M_{\odot}$. Reducing the bubble mass by

excluding extra gas along the line of sight brings the estimates closer in line with both empirical and theoretical models for feedback.

Table 2 lists the physical parameters of all of the Taurus bubbles. It includes the estimates from Li et al. (2015) and our models ME1 and MF. Li et al. (2015) do not consider the correction factors for bubble mass, momentum, and energy, due to the limited ^{13}CO velocity range, so to make a fair comparison, we do not apply correction factors to the predictions from models ME1 and MF.

3.3. Assessing the Global Impact of Feedback: Full Taurus Map

3.3.1. Feedback Features Identified in the Full Map

We apply the CASI-3D models to the complete Taurus map to predict all of the emission associated with feedback. We divide the Taurus map into smaller cubes as discussed in Section 2.4. To create the full prediction map, we adopt the largest value from the overlapping predictions at each pixel. Note that the 5×5 pixel regions in the map corners have only one cube prediction for each pixel.

To check the accuracy of this method, we compare the model predictions of the postage stamps and those from the large map. Figure 16 shows that the large map prediction captures the bubble rims better than the single postage stamp predictions.

Figures 17 and 18 show the predictions from models ME1 and MF for the whole Taurus map. The CASI-3D model predictions cover almost all of the previously identified bubble regions and predict additional feedback regions in the Taurus map. The new predictions are correlated with the locations of Class III YSOs, as shown on the map. Figure 17 shows that most predictions are close to several groups of YSOs. For example, new bubble N3, which was not previously identified, seems to enclose a large group of YSOs. This suggests that the YSOs are shaping the surrounding clouds through their feedback and creating a wind signature in the ^{13}CO spectra. We further discuss the newly detected feedback regions below.

We identify three types of bubbles in our model predictions: high-confidence bubbles that were identified by the previous observational survey (red boxes), high-confidence bubbles that we believe are real bubbles that were missed in the previous survey (yellow boxes), and low-confidence bubbles that are new bubbles found by our models but we believe are less certain (white boxes). The first category of high-confidence bubbles corresponds to “true positives.” The second category of missing high-confidence bubbles corresponds to “true negatives,” and the final category of low-confidence bubbles may represent “false positives.”

First, we discuss the high-confidence bubbles (true positives) that are consistent with the previous human identifications. These bubbles have a clear ring or arc-like structure and have at least one YSO inside. Bubbles H1, H2, H3, and H4 correspond to TMB 37, TMB 29, TMB12, and TMB7 in Li et al. (2015), respectively. These bubbles are identified by both model ME1 and MF, although the extent of the emission in model MF may be smaller if the fraction of the mass coming from feedback is predicted to be low.

Next, we discuss the high-confidence bubbles that were not included in Li et al. (2015). These bubbles have a clear bubble rim morphology and have YSOs nearby if not directly within the bubble center. For example, in the yellow box N1 we see

Table 2
Physical Parameters of Taurus Bubbles

Bubble ID	Rank	Li+ (2015)			ME1			MF		
		Mass ($M_{\odot}$)	Momentum ($M_{\odot} \text{ km s}^{-1}$)	Energy (10^{44} erg)	Mass ($M_{\odot}$)	Momentum ($M_{\odot} \text{ km s}^{-1}$)	Energy (10^{44} erg)	Mass ($M_{\odot}$)	Momentum ($M_{\odot} \text{ km s}^{-1}$)	Energy (10^{44} erg)
TMS_1	B	59	65	7	15	39	11	1.9	6.2	2.25
TMS_2	B	31	34	4	43	32	4	9.1	8.2	1.29
TMS_3	B	129	207	33	120	169	35	10.7	15.1	3.33
TMS_4	B	56	62	7	42	65	14	3.6	5.4	1.14
TMS_5	B	32	52	8	56	45	5	3.6	3.5	0.54
TMS_6	B	33	37	4	10	9	1	2.7	3.2	0.64
TMS_7	C	91	191	40	95	73	9	10.5	9.7	1.50
TMS_8	B	46	97	20	33	48	12	2.8	4.5	1.12
TMS_9	C	22	17	1	39	34	4	5.3	4.3	0.58
TMS_10	C	217	282	37	228	272	57	24.8	35.8	8.26
TMS_11	A	78	149	28	94	70	9	7.1	5.9	0.81
TMS_12	A	45	108	26	57	43	5	3.7	3.0	0.39
TMS_13	B	84	176	37	112	88	11	8.1	8.6	1.46
TMS_14	C	18	33	6	13	44	17	2.8	9.0	3.71
TMS_15	A	98	107	12	109	90	12	6.0	7.6	1.68
TMS_16	A	57	92	15	75	60	8	1.5	2.4	0.64
TMS_17	C	10	8	1	15	19	3	1.4	2.4	0.67
TMS_18	C	54	103	19	72	63	9	2.5	3.0	0.57
TMS_19	A	41	78	15	66	77	13	3.5	4.0	0.68
TMS_20	B	12	13	1	21	28	5	1.8	2.3	0.42
TMS_21	B	11	24	5	3	15	9	1.1	4.6	2.38
TMS_22	B	119	131	14	134	168	33	12.3	18.4	3.92
TMS_23	C	10	11	1	24	21	3	0.1	0.1	0.02
TMS_24	C	20	22	2	29	52	10	1.8	2.9	0.56
TMS_25	C	9	26	8	7	11	2	0.2	0.6	0.22
TMS_26	B	11	23	5	8	29	12	3.0	10.6	4.71
TMS_27	C	46	111	26	24	82	33	9.4	32.3	13.87
TMS_28	C	205	656	209	293	365	68	14.2	18.4	3.50
TMS_29	A	91	119	15	120	145	27	7.0	7.3	1.19
TMS_30	C	420	672	107	576	636	106	23.2	26.3	4.65
TMS_31	C	62	81	10	43	89	26	6.6	23.9	10.33
TMS_32	C	12	22	4	7	29	14	2.6	12.6	6.74
TMS_33	B	92	74	6	143	153	25	5.6	5.9	0.96
TMS_34	C	67	182	49	36	39	7	4.7	5.3	1.51
TMS_35	B	30	39	5	22	27	4	4.4	3.8	0.64
TMS_36	A	29	47	7	16	25	6	3.7	4.2	1.00
TMS_37	B	25	53	11	19	40	12	7.0	17.6	6.96

the bubble rim and the cavity. Moreover, one Class III YSO is centered in the cavity, which is likely to be the driving source of the bubble. The predictions from both ME1 and MF for N1 highlight the bubble rim. In another example, N3, we can easily identify the bubble rim in Figure 17. Supporting the bubble's existence is a group of Class III YSOs inside its rim. However, when we look at the prediction from model MF for N3 in Figure 18, we cannot see the bubble rim prediction. This suggests that there is likely a small amount of mass coming from the feedback.

Finally, we discuss the low-confidence bubbles. These bubbles tend not to be associated with any YSOs. In addition to the Class III YSOs identified in Kraus et al. (2017), we check all types of YSOs in Taurus that were identified by Rebull et al. (2010). These are plotted in Figures 31 and 32 in Appendix E. We highlight four such bubbles in white boxes in Figure 17. We note that a number of the bubbles identified by Li et al. (2015) do not contain any Class III YSOs, such as L4. In these cases, the driving source may have moved out of the bubble, or the YSO census may be incomplete. Another possible

explanation is that although the morphology is circular or arc-like, they are caused by cloud turbulence, which causes coherent motion across several velocity channels. It is difficult to distinguish the bubble structure from turbulent patterns when a circular or arc-like pattern appears across multiple channels. However, we believe this last explanation is unlikely, because we include pure turbulence snapshots in the training set as negative training images. Thus, CASI-3D should not be prone to misidentifying turbulent patterns as bubbles.

Overall, we conclude that the two CNN models perform as well as or better than “by-eye” visual identifications of bubbles. They appear to reasonably predict both the bubble position and the fraction of mass coming from feedback.

3.3.2. Mass, Momentum, and Energy of the Feedback Identified in the Full Taurus Cloud

We now calculate the feedback mass, momentum, and energy in Taurus based on the predictions from models ME1 and MF. Table 3 lists the feedback properties calculated in this work and those calculated in Li et al. (2015).




Figure 16. Comparison of the prediction of a postage stamp and that from the large map on bubble TMB29. Left: integrated ^{13}CO intensity. The red stars indicate Class III YSOs from the Kraus et al. (2017) catalog. Middle: integrated prediction from model ME1 run on the postage stamp shown in the left panel. Right: integrated ME1 prediction from the full map prediction, reconstructed from overlapping postage stamps.

Model ME1 predicts $2630 M_{\odot}$ of gas associated with feedback, which is consistent within a factor of 2 with the feedback mass calculated in Li et al. (2015). However, model MF predicts that only $275 M_{\odot}$ of gas is associated with feedback, which is an order of magnitude smaller than the previous calculations. The smaller amount of feedback mass is also consistent with the total stellar mass in Taurus, which is estimated to be on the order of $200 M_{\odot}$ in Kraus et al. (2017). The feedback mass predicted by model ME1 and that calculated from Li et al. (2015) are 10 times the stellar mass, which is inconsistent with the expected amount of gas entrained by feedback (Offner & Chaban 2017).

Despite the detailed machine-learning identification, we must still confront the challenge of how to disentangle feedback from the bulk cloud motion. For example, Taurus has a velocity gradient that stretches from the southeast to northwest. Not accounting for this gradient may artificially enhance the feedback total. The most accurate way to account for the bulk motion is not clear; thus, we present two approaches to calculating the feedback momentum and energy. The first way treats the molecular cloud as a whole, with the same fixed central velocity. We shift the central velocity to zero and calculate the 1D momentum and 1D energy channel by channel as described in Section 3.1. The second approach is similar but treats the molecular cloud locally, which means the cloud does not have a fixed central velocity but has a central velocity gradient across the entire cloud. We subtract the central velocity pixel by pixel and then calculate the momentum and energy channel by channel. To convert the 1D line-of-sight estimates to 3D, we make the assumption of isotropic expansion to calculate the 3D momentum and 3D energy. Finally, in Section 3.1, we assessed the model accuracy using synthetic observations and found that 20% of the momentum and 44% of the energy are missed because of the limited CO velocity coverage. Considering the limited velocity range of the ^{13}CO data cube, we apply these correction factors for the missing momentum and energy here.

The momentum estimate without the velocity gradient treatment from model ME1 is close to that from Li et al. (2015). The momentum estimate with the velocity gradient treatment from model ME1 is 38% smaller than the calculation in Li et al. (2015). Once corrected for the extra ^{13}CO emission from the foreground or background, the momentum (with and

without the velocity gradient treatment) predicted by model MF is an order of magnitude smaller than the calculation in Li et al. (2015).

Both energy estimates (with and without the velocity gradient treatment) from model ME1 are within a factor of 2 compared to the energy calculated in Li et al. (2015). In contrast, model MF implicitly corrects for the extra ^{13}CO emission coming from the foreground or background, such that the predicted energy is an order of magnitude smaller than that calculated in Li et al. (2015). We discuss the implications in the following section.

3.3.3. Assessing the Relative Energies of Turbulence and Feedback

In this section, we compare the total energy associated with feedback and the total cloud turbulent energy. The relative magnitude of these energies impacts the cloud lifetime and whether turbulence can slow collapse by providing pressure support against self-gravity. Often, the impact of feedback is weighed against the rate of turbulence dissipation. We follow Li et al. (2015) and define the turbulent dissipation rate as

$$L_{\text{turb}} = \frac{E_{\text{turb}}}{t_{\text{diss}}}, \quad (4)$$

where t_{diss} is the turbulent dissipation time. The method to estimate the turbulent dissipation time in Li et al. (2015) is from Mac Low (1999),

$$t_{\text{diss}} \sim \left(\frac{0.39\kappa}{\mathcal{M}_{\text{rms}}} \right) t_{\text{ff}}, \quad (5)$$

where t_{ff} is the freefall timescale, $\mathcal{M}_{\text{rms}}$ is the Mach number of the turbulence, and κ is the ratio of the driving length to the Jean's length of the cloud. For $\mathcal{M}_{\text{rms}} = 5$ and a freefall time $t_{\text{ff}} = 7 \times 10^6$ yr, which assumes a mean cloud number density of $n = 20 \text{ cm}^{-3}$, the turbulent dissipation rate is $3.1 \times 10^{33} \text{ erg s}^{-1}$. However, Taurus is not a uniform sphere, and the mean number density of $n = 20 \text{ cm}^{-3}$ adopted by Li et al. (2015) is too low. The typical mean number density of a molecular cloud is around $n = 100 \text{ cm}^{-3}$, which gives $t_{\text{ff}} = 3.3 \times 10^6$ yr and $L_{\text{turb}} = 6.8 \times 10^{33} \text{ erg s}^{-1}$. Arce et al. (2010) and Narayanan et al. (2012) also adopt this method to




Figure 17. The ^{13}CO integrated intensity of the Taurus molecular cloud overlaid with the integrated prediction of feedback position from ME1 along velocity channels in red. The arcs in yellow indicate the position of previously identified bubbles in Li et al. (2015). The star symbols demonstrate the locations of the Class III YSOs from Kraus et al. (2017).




Figure 18. Same as Figure 17 but predicted by MF.

calculate the turbulent dissipation rates in Perseus and Taurus, respectively.

One caveat here is that the equation to calculate the turbulent dissipation rate is obtained from simulations, which depend on the initial conditions and the way turbulence is driven.

The energy injection rate is defined as $L_{\text{bubble}} = E_{\text{bubble}}/t_{\text{kinetic}}$, where E_{bubble} is the kinetic energy of the bubble and t_{kinetic} is the kinetic timescale of the bubble. The kinetic timescale of the bubble can be calculated as $t_{\text{kinetic}} = R/V_{\text{exp}}$, where R is the

radius of the bubble and V_{exp} is the expansion velocity of the bubble. We find that the energy injection rate from bubbles in ME1 is $L_{\text{turb,ME1}} = 1.0 \times 10^{34} \text{ erg s}^{-1}$, which is slightly larger than the turbulent dissipation rate of the cloud. If we subtract the velocity gradient, the energy injection rate from bubbles is $L_{\text{turb,ME1,G}} = 2.8 \times 10^{33} \text{ erg s}^{-1}$, which is about half of the turbulent dissipation rate of the cloud. In summary, like Li et al. (2015), we conclude that feedback is sufficient to maintain the current level of cloud turbulence.

Table 3
Properties of Feedback in the Taurus Molecular Cloud^a

Model	Without Subtracting the Velocity Gradient				Subtracting the Velocity Gradient			
	M^b ($M_\odot$)	P_{3D} ($M_\odot \text{ km s}^{-1}$)	E_{3D}^c ($\times 10^{46} \text{ erg}$)	$\dot{E}^d$ ($\times 10^{33} \text{ erg s}^{-1}$)	M^b ($M_\odot$)	P_{3D} ($M_\odot \text{ km s}^{-1}$)	E_{3D}^c ($\times 10^{46} \text{ erg}$)	$\dot{E}^d$ ($\times 10^{33} \text{ erg s}^{-1}$)
Li+ (2015)	1707 (11.4%)	3780	9.2 (28.8%)	6.4 (94.1%)	...	...	...	...
ME1	2894 (19.3%)	4366	15 (46.2%)	10 (153%)	2894 (19.3%)	2339	4.0 (12.6%)	2.8 (40.9%)
MF	302 (2.0%)	609	2.8 (8.6%)	2.0 (28.6%)	302 (2.0%)	366	0.96 (3.0%)	0.67 (9.8%)

Notes.

^a Model name, feedback bubble mass, 3D feedback momentum, 3D feedback energy, and energy injection rate from feedback bubbles. The numbers in the table consider the correction factors that are due to the limited velocity range of the ^{13}CO data cube.

^b The number in parentheses indicates the percentage of feedback mass compared to the whole molecular cloud mass (Pineda et al. 2010).

^c The number in parentheses indicates the percentage of feedback energy compared to the whole molecular cloud turbulent energy (Li et al. 2015).

^d The number in parentheses indicates the percentage of energy injection rate from feedback bubbles compared to the turbulent dissipation rate of the cloud. The turbulent dissipation rate adopted here is $L_{\text{turb}} = 6.8 \times 10^{33} \text{ erg s}^{-1}$, which assumes a mean cloud density of $n = 100 \text{ cm}^{-3}$. This turbulent dissipation rate is about two times higher than that from Li et al. (2015), which assumes a lower mean cloud density of $n = 20 \text{ cm}^{-3}$.

However, we have shown that model ME1 overestimates the energy because excess foreground and background material is included in the calculation. Consequently, we find that after recalculating the feedback energy using the more accurate model MF prediction, the kinetic energy from the feedback decreases by an order of magnitude, which means the energy injection rate from stars is smaller by an order of magnitude: $L_{\text{turb, MF}} = 2.0 \times 10^{33} \text{ erg s}^{-1}$. Under this circumstance, the energy injection rate from feedback is 29% of the turbulent dissipation rate of the cloud. If we subtract the velocity gradient, the energy injection rate from feedback is $L_{\text{turb, MF, G}} = 6.7 \times 10^{32} \text{ erg s}^{-1}$, which is an order of magnitude smaller than the turbulent dissipation rate. This indicates that some additional energy is needed to drive turbulence in the Taurus molecular cloud, which could be provided by outflows, for example. Feedback from bubbles may not be sufficient to maintain the cloud turbulence over long timescales.

The Taurus molecular cloud is host to an older population of stars ($\tau \sim 10\text{--}20 \text{ Myr}$), which indicates the lifetime of the cloud is at least 10–20 million years (Kraus et al. 2017). However, this lifetime is much longer than the gravitational collapse freefall time of the Taurus molecular cloud estimated from ^{12}CO , which is 3.3 million years. This suggests that there must be energy injected to support the cloud against gravitational collapse, which suggests feedback is playing some role in driving turbulence but is not dominant.

3.3.4. Quantifying the Impact of Feedback with Turbulent Statistics

With an accurate prediction of the position of feedback in hand, we compute multiple astrostatistics to study the different properties between regions with and without feedback in Taurus. We adopt the statistical analysis package TURBUSTAT to conduct the statistical analysis (Koch et al. 2017, 2019). TURBUSTAT contains 15 different statistics, but here we consider only the spatial power spectrum (SPS) and the covariance matrix used to compute the principal component analysis (PCA). We adopt these statistics because they have previously been shown to be sensitive to feedback, as discussed in the introduction. The SPS is defined as the square of the 2D Fourier transform of an image:

$$\begin{aligned} \mathcal{P}(k) &= \sum_{|\vec{k}|=k} |\mathcal{M}_0(\mathbf{k})|^2 \\ &= \left| \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} M_0(\mathbf{x}) e^{-2\pi i \mathbf{k} \cdot \mathbf{x}} d\mathbf{x} \right|^2. \end{aligned} \quad (6)$$

It is applied to the integrated intensity map. The covariance matrix is defined as

$$C_{jk} = \frac{1}{n} \sum_{i=1}^n X_{ij} X_{ik}, \quad (7)$$

where

$$X_{ij} = T(r_i, v_j) - \left[\sum_{k=1}^n T(r_k, v_j) \right] / n, \quad (8)$$

in which $T(r_i, v_j)$ is the spectral cube, where $r_i = (x_i, y_i)$ is the position on the sky and v_j indicates the spectral velocity channel. It provides information about velocity correlations. In addition to these two statistics, we also consider the distribution of line widths of the feedback and nonfeedback gas, as well as the distance between YSOs and pixels associated with feedback.

Figure 19 shows the SPS of the full ^{13}CO integrated intensity map of Taurus, the SPS of the region where the emission is above 0.2 K (i.e., excluding noise), and the SPS of the model ME1 and MF predicted feedback regions. Figure 19 shows that the slope of the SPS is flattened over the feedback injection region. If the emission is optically thin and the temperature is roughly constant, this indicates mass or energy has been injected into smaller scales by the feedback. Here, the ^{13}CO is mostly optically thin, with the exception of dense cores.

Next, in Figure 20 we present the covariance matrices of the velocity channels for the full ^{13}CO integrated intensity map of Taurus, the high signal-to-noise ratio region, and the prediction of the models ME1 and MF. For comparison, Figure 20 also shows the covariance matrices calculated using the synthetic data. The covariance matrices of the predicted feedback regions clearly show off-diagonal velocity features, which indicate coherent motions at these velocities. These features can be characteristic of the expansion of bubbles or high-velocity gas (Boyden et al. 2016), but they may also represent coherent cloud motions (e.g., Feddersen et al. 2019). In either case, the clear differences between the identified feedback and nonfeedback gas underscore that CASI-3D is indeed identifying statistically distinct regions.

Next, we assess the relative distance to the YSO locations, which provide additional evidence that our regions are associated with feedback. Figure 21 shows the distribution of the projected distances between the YSOs and the emitting gas, and the distribution of the projected distances between YSOs and the feedback gas predicted by ME1 and MF. The median value of




Figure 19. SPS of the full ^{13}CO integrated intensity map of Taurus, the SPS of the emission regions (excluding noise regions) where the emission is above 0.2 K, and the SPS of the ME1 and MF predicted feedback regions.

the projected distance between the YSOs and the feedback gas is closer than that between the YSOs and all of the emitting gas. The typical distance between the YSOs and the feedback gas is 0.7 pc, which is also the typical size of the bubbles.

Finally, we expect feedback regions to have larger velocity dispersions. Figure 22 shows the distribution of the FWHM of the high signal-to-noise ratio emission region and the FWHM of the ME1 and MF predicted feedback regions. The median values of the FWHM of the feedback regions are indeed larger than that of the FWHM of the full map. The higher FWHM indicates larger velocities in the spectrum associated with feedback.

4. Conclusions

We adopt a deep learning method, CASI, and extend it to 3D (CASI-3D) to identify stellar feedback features in 3D CO spectral cubes. By creating different training sets, we develop two deep machine-learning tasks. Task I predicts the position of feedback. Task II predicts the fraction of the mass coming from feedback. Our main findings are the following:

1. CASI-3D is a powerful method for identifying bubbles. CASI-3D performs well on synthetic test data and recovers feedback with an accuracy of 4% on a pixel level.
2. CASI-3D successfully predicts hidden information, such as the fraction of mass coming from feedback.

3. We apply CASI-3D to the ^{13}CO observations of the Taurus molecular cloud and show that CASI-3D successfully identifies previously known, visually identified bubbles.
4. We find that training Task I reproduces the mass, momentum, and energy of individual bubbles inferred by human visual identifications. In contrast, Task II, which is trained on the feedback mass fraction, indicates that the true mass, momentum, and energy are an order of magnitude lower.
5. CASI-3D suggests that previous studies overestimate feedback mass and energy in the Taurus molecular cloud. The feedback mass is overestimated by a factor of 5. The feedback energy is overestimated by a factor of 5 compared to that calculated without subtracting the velocity gradient over the full map, and it is overestimated by a factor of 10 compared to that calculated with subtracting the velocity gradient over the full map.
6. We carry out an analysis of the SPS to quantify the turbulence properties in the feedback and nonfeedback regions. We show that feedback flattens the slope of the SPS of the full ^{13}CO integrated intensity map of Taurus, indicating that mass or energy has been injected at smaller scales by feedback.
7. We calculate the covariance matrix and show that the presence of feedback appears as off-diagonal peaks in the covariance matrices.




Figure 20. Covariance matrices of the velocity channels on the full ^{13}CO integrated intensity map of Taurus, the covariance matrices of the emission regions, and the covariance matrices of the ME1 and MF predicted feedback regions. Left panel: covariance matrices of the velocity channels on synthetic data. Right panel: covariance matrices of the velocity channels on Taurus ^{13}CO data.




Figure 21. Distribution of the projected distance between YSOs and the emitting gas, and the distribution of the projected distance between YSOs and the feedback gas predicted by ME1 and MF.

8. The median value of the projected distance between YSOs and the feedback gas (0.64 pc predicted by model ME1 and 0.75 pc predicted by model MF) is closer than that between YSOs and all of the emitting gas (0.84 pc). The median value of the FWHM of the feedback regions (1.2 km s^{-1} predicted by model ME1 and 1.1 km s^{-1} predicted by model MF) is larger than that of the FWHM of the full emitting regions (0.9 km s^{-1}).

In future work, we plan to apply CASI-3D to other star-forming regions and other types of feedback, such as protostellar outflows (Arce et al. 2010).

D.X., S.S.R.O., R.A.G., and C.V.O. were supported by NSF grant AST-1812747. S.S.R.O. also acknowledges support from NSF Career grant AST-1650486. The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this paper.




Figure 22. Distribution of the FWHM of the emission regions, and the FWHM of the ME1 and MF predicted feedback regions.

Appendix A CASI-3D Parameters

A.1. Down-sampling Methods

We test two widely used down-sampling methods to reduce the size of the data: max pooling and average pooling. Max pooling picks out the largest value to replace its adjacent pixels. Max pooling can extract the most important features, but it is not proficient in dealing with different noise backgrounds. Since all large-map sky surveys are conducted through substantial observing periods with different weather conditions and with different baselines, the noise level is different in different patches of the large map. When applying max pooling to down-sample the data, the boundary between patches distinctly appears, which makes the data inconsistent across the map. On the other hand, average pooling extracts features smoothly, and it preserves the overall value during down-sampling. Figure 23 shows an




Figure 23. Comparison of two different down-sampling methods tested on ^{13}CO Taurus molecular cloud data.



Figure 24. Performance of model MF adopting different loss functions to predict the fraction of the mass that comes from stellar feedback on a test bubble. Left: integrated ^{13}CO intensity. Second from left: integrated true feedback fraction. Right panels: models using the IoU, MSE, and IoU+MSE loss functions, respectively.

example of the two different down-sampling methods tested on ^{13}CO Taurus molecular cloud data.

A.2. Loss Function

We test three types of loss functions—MSE, IoU, and a combination of MSE and IoU—to predict the fraction of the mass that comes from stellar feedback. Figure 24 shows the performance of the model using different loss functions on a test bubble. The model adopting IoU as the loss function can capture the morphology of the bubble clearly but misses the value information. The IoU model predicts almost unity at the feedback position but does not reflect the actual fraction value. The MSE model is able to capture the position of larger feedback values but underestimates the smaller values, which is useful for predicting the emission but not the fraction. The model adopting both MSE and IoU as the loss function performs the best. This model not only captures the distinct bubble morphology but also returns reasonable fraction values.

Appendix B Training Sets

A.1. Comparison of ^{13}CO Emission with Different Cloud Thicknesses

Figure 25 shows the difference in the synthetic observations between the whole cube and the cropped data for ^{13}CO . The ^{13}CO bubble rim is embedded in the diffuse gas emission, and the bubble cavity is not clear in the integrated intensity map when the thickness of the cloud is 5 pc. Because ^{12}CO is even

more optically thick, ^{12}CO is not an appropriate proxy to trace stellar feedback winds (e.g., Arce et al. 2011; Li et al. 2015). When the thickness of the cloud becomes smaller, the bubble rim and its cavity are recognizable in the ^{13}CO integrated intensity map. Although some bubble rims or their cavities are not distinct in the integrated intensity map of ^{13}CO , these feedback features become recognizable in PPV space.

A.2. Different Definitions for the Bubble Extents

In this section, we assess the impact of different choices for the bubble definition on the results. In addition to the bubble definition described in Section 2.3.2, we also examine the tracer field in the simulation data. The gas adjacent to the tracer gas has a velocity vector going outward, which indicates the feedback gas compresses the ambient gas without direct contact. Although the fraction of feedback gas compared to the entire amount of gas contained in these voxels is almost zero, the adjacent layer contributes to the momentum and the energy of the cloud. Consequently, we define the tracer field with the velocity vector going outward from the central stars, which increases the mass of the feedback bubble by a factor of 3. We furthermore test a temperature cut at $T \geq 12$ K near the tracer gas to calculate the bubble mass. The simulation data cubes have an average temperature of 10 K. The temperature drops quickly from the bubble rim to the ambient gas. We compare the different tracer definitions in Figure 26. Both the velocity cut plus the tracer field and the temperature cut plus the tracer field are slightly larger in area than the original tracer field. Both of these definitions yield bubbles that are similar in




Figure 25. Integrated intensity of ^{13}CO ($J = 1 - 0$). Upper left: integrated intensity of ^{13}CO generated using the whole data cube. Upper middle and upper right: integrated intensity of ^{13}CO generated using the cropped data cube. Bottom left: integrated intensity of ^{13}CO wind tracer generated using the whole data cube. Bottom middle and bottom right: integrated intensity of ^{13}CO wind tracer generated using the cropped data cube. “D” is the line-of-sight thickness of the data cube.




Figure 26. Different definitions for a synthetic bubble. First panel: the ^{13}CO integrated intensity map. Second panel: the original definition of the bubble using the ORION tracer field integrated over the velocity channels. Third panel: the temperature-based definition ($T > 12$ K plus the tracer field) integrated over the velocity channels. Fourth panel: the velocity-based definition (gas with expanding velocities plus the tracer field) integrated over the velocity channels.

shape. Because there are five individual stars in the simulation box, the bubbles generated by these stars are easily connected to each other during the expansion. This affects the gas velocities, which makes it difficult for us to define the gas flow direction and determine which expansion is part of the shell. Under this circumstance, the temperature cut is a better option for defining the bubble boundary. We compare the bubble mass calculated from the velocity-based bubble definition and the temperature-based bubble definition in Figure 27. Larger bubbles are more likely to overlap during the expansion, which makes the velocity-based bubble definition mass slightly

smaller than the the temperature-based bubble definition mass. Overall, we conclude the temperature-based bubble definition is the most appropriate definition of the bubble boundary.

A.3. Comparison of the Training and Observed Bubble Mass Distributions

In this section, we examine the distribution of the bubble masses in the training set. Figure 28 shows that the maximum bubble mass in the training set is $\sim 120 M_{\odot}$, and it spans the range of the bubble masses in the test samples in Section 3.1.




Figure 27. Comparison between the velocity-cut tracer mass and the temperature-cut tracer mass.




Figure 28. Distribution of bubble masses in the training set.

Moreover, to extend the range of bubble masses, we have included “zoomed-in” synthetic observations. In these 64×64 postage stamps, the original bubble is enlarged by a factor of 2 in both length and width, which indicates the bubble area and the mass both increase by a factor of 4. Because CASI-3D takes postage stamp cubes as inputs, regardless of the actual physical size of the cubes, the training set spans bubble masses up to $\sim 4 \times 120 M_{\odot}$, and it spans the range of the individual bubble masses in observations in Section 3.2. In some cases, a single 64×64 postage stamp cannot cover an entire bubble. Only part of the bubble appears within the input window, such as an arch or a half circle. These cases in the training set are consistent with the cases of larger bubbles that are contained in the full map prediction in Taurus. Thus, to obtain the masses of the largest bubbles in Taurus, we combine a stack of postage stamps that cover different parts of each bubble to get the full prediction and then calculate the bubble mass as described in Section 3.2.

Appendix C

Excitation Temperature Selection and Impact

In this section, we explore the uncertainty in the bubble masses that is due to the choice of excitation temperature. We find that




Figure 29. Ratio between the mass estimated from ^{13}CO assuming a 25 K excitation temperature and the true mass calculated from the simulations.

25 K (e.g., Narayanan et al. 2012; Li et al. 2015) is the most appropriate choice for converting ^{13}CO emission to column density in the synthetic observations. We show the ratio between the mass estimated from ^{13}CO assuming a 25 K excitation temperature and the true mass calculated from the simulations in Figure 29. The ratio is within a factor of 2 of unity when assuming LTE and a 25 K excitation temperature, which in turn demonstrates that both LTE and the choice of 25 K are reasonable.

Under the assumption of LTE, the mass estimation goes linearly with the excitation temperature. From previous feedback mass estimates (e.g., Arce et al. 2011; Li et al. 2015), the choice of excitation temperature ranges from 10 to 50 K. This could introduce a factor of 2 uncertainty in the mass estimation, but it cannot account for a factor of 10.

Appendix D

Assessing the Sensitivity of the Data Window

In this section, we check the sensitivity of the data window to feedback as a function of voxel location. As discussed in Section 2.2.1, the CASI-3D models predict the full Taurus map feedback using a stack of $64 \times 64 \times 32$ cubes, in total 11,340 cubes. We examine the “response” of each voxel in a $64 \times 64 \times 32$ cube. We define the response as the fraction of the stacked voxels over the 11,340 cubes that are detections in a $64 \times 64 \times 32$ cube. A detection is defined as a voxel above 90% of the maximum prediction value for all overlapping voxels at the corresponding full map location. Figure 30 shows the response integrated over the velocity channels. The central region of the postage stamp is the most sensitive region, where a higher fraction of the stacked voxels fall above the prediction threshold. The boundary region of the postage stamp is less sensitive to features, and detection is less efficient. This figure illustrates that choosing the appropriate cube offset is important to achieving the best sensitivity. To ensure that all data are covered by the highly sensitive part of the window, the maximum step size should not exceed 16 pixels. In our stacked prediction, we adopt a step size of five pixels, which is smaller than the required step size, to crop the full Taurus map.




Figure 30. Integrated response of a voxel in the cube, that is, the fraction of voxels predicted to be associated with feedback in the same position in the stack of cubes summed over all velocity channels.

We note that the range of the window sensitivity likely depends on the training data and the target feature size. Here, we aim to find bubbles that have typical sizes greater than ~ 16 pixels or a quarter of the cube length. We recommend that other users of CASI-3D check the window sensitivity for their problem to determine the appropriate offset size when applying CASI-3D to large data maps.

Appendix E

YSOs in the Taurus Molecular Cloud

Figures 31 and 32 show the ^{13}CO integrated intensity of Taurus overlaid with the integrated prediction of feedback from models ME1 and MF (red). The arcs in yellow indicate the position of the previously identified bubbles from Li et al. (2015). The red stars indicate the locations of the Class III YSOs from Kraus et al. (2017). The white stars represent the locations of the YSOs from Rebull et al. (2011). The white cross symbols indicate the locations of the YSOs from Rebull et al. (2010).




Figure 31. The ^{13}CO integrated intensity of Taurus overlaid with the integrated prediction of feedback from models ME1 and MF (red). The arcs in yellow indicate the position of the previously identified bubbles from Li et al. (2015). The red stars indicate the locations of the Class III YSOs from Kraus et al. (2017). The white stars represent the locations of the YSOs from Rebull et al. (2011). The white cross symbols indicate the locations of the YSOs from Rebull et al. (2010).




Figure 32. Same as Figure 31 but predicted by model MF.

ORCID iDs

Stella S. R. Offner <https://orcid.org/0000-0003-1252-9916>

Robert Gutermuth <https://orcid.org/0000-0002-6447-899X>

References

- Arce, H. G., Borkin, M. A., Goodman, A. A., et al. 2010, *ApJ*, **715**, 1170
- Arce, H. G., Borkin, M. A., Goodman, A. A., Pineda, J. E., & Beaumont, C. N. 2011, *ApJ*, **742**, 105
- Bally, J. 2016, *ARA&A*, **54**, 491
- Beaumont, C. N., Goodman, A. A., Kendrew, S., Williams, J. P., & Simpson, R. 2014, *ApJS*, **214**, 3
- Beaumont, C. N., Williams, J. P., & Goodman, A. A. 2011, *ApJ*, **741**, 14
- Boyden, R. D., Koch, E. W., Rosolowsky, E. W., et al. 2016, *ApJ*, **833**, 233
- Churchwell, E., Povich, M. S., Allen, D., et al. 2006, *ApJ*, **649**, 759
- Churchwell, E., Watson, D. F., Povich, M. S., et al. 2007, *ApJ*, **670**, 428
- Dullemond, C. P., Juhasz, A., Pohl, A., et al. 2012, RADMC-3D: a multi-purpose radiative transfer tool, Astrophysics Source Code Library, ascl:1202.015
- Feddersen, J. R., Arce, H. G., Kong, S., et al. 2019, *ApJ*, **875**, 162
- Frank, A., Ray, T. P., Cabrit, S., et al. 2014, in *Protostars and Planets VI*, ed. H. Beuther et al. (Tucson, AZ: Univ. Arizona Press), 451
- Hartigan, P., Edwards, S., & Ghandour, L. 1995, *ApJ*, **452**, 736
- He, K., Zhang, X., Ren, S., & Sun, J. 2016, in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition* (Piscataway, NJ: IEEE), 770
- Hollenbach, D. J., & Tielens, A. G. G. M. 1999, *RvMP*, **71**, 173
- Jayasinghe, T., Dixon, D., Povich, M. S., et al. 2019, *MNRAS*, **488**, 1141
- Koch, E. W., Rosolowsky, E. W., Boyden, R. D., et al. 2019, *AJ*, **158**, 1
- Koch, E. W., Ward, C. G., Offner, S., et al. 2017, *MNRAS*, **471**, 1506
- Kraus, A. L., Herczeg, G. J., Rizzuto, A. C., et al. 2017, *ApJ*, **838**, 150
- Li, H., Li, D., Qian, L., et al. 2015, *ApJS*, **219**, 20
- Li, P. S., Martin, D. F., Klein, R. I., et al. 2012, *ApJ*, **745**, 139
- Mac Low, M.-M. 1999, *ApJ*, **524**, 169
- Molinari, S., Swinyard, B., Bally, J., et al. 2010, *PASP*, **122**, 314
- Nakamura, F., & Li, Z.-Y. 2008, *ApJ*, **687**, 354
- Nakamura, F., Miura, T., Kitamura, Y., et al. 2012, *ApJ*, **746**, 25
- Narayanan, G., Heyer, M. H., Brunt, C., et al. 2008, *ApJS*, **177**, 341
- Narayanan, G., Snell, R., & Bemis, A. 2012, *MNRAS*, **425**, 2641
- Ntampaka, M., Zuhone, J., Eisenstein, D., et al. 2019, *ApJ*, **876**, 82
- Offner, S. S. R., & Arce, H. G. 2015, *ApJ*, **811**, 146
- Offner, S. S. R., & Chaban, J. 2017, *ApJ*, **847**, 104
- Offner, S. S. R., & Liu, Y. 2018, *NatAs*, **2**, 896
- Peek, J. E. G., Heiles, C., Douglas, K. A., et al. 2011, *ApJS*, **194**, 20
- Pineda, J. L., Goldsmith, P. F., Chapman, N., et al. 2010, *ApJ*, **721**, 686
- Qian, L., Li, D., Offner, S., et al. 2015, *ApJ*, **811**, 71
- Quillen, A. C., Thorndike, S. L., Cunningham, A., et al. 2005, *ApJ*, **632**, 941
- Rebull, L. M., Koenig, X. P., Padgett, D. L., et al. 2011, *ApJS*, **196**, 4
- Rebull, L. M., Padgett, D. L., McCabe, C.-E., et al. 2010, *ApJS*, **186**, 259
- Ridge, N. A., Di Francesco, J., Kirk, H., et al. 2006, *AJ*, **131**, 2921
- Ronneberger, O., Fischer, P., & Brox, T. 2015, arXiv:1505.04597
- Shu, F., Najita, J., Ostriker, E., et al. 1994, *ApJ*, **429**, 781
- Simpson, R. J., Povich, M. S., Kendrew, S., et al. 2012, *MNRAS*, **424**, 2442
- Smith, N. 2014, *ARA&A*, **52**, 487
- Van Oort, C. M., Xu, D., Offner, S. S. R., et al. 2019, *ApJ*, **880**, 83
- Veit, A., Wilber, M. J., & Belongie, S. 2016, in *Advances in Neural Information Processing Systems*, ed. D. D. Lee et al. (Red Hook, NY: Curran Associates, Inc.), 550
- Xu, D., & Offner, S. S. R. 2017, *ApJ*, **851**, 149