Wall Interactions with AMB+ Model Particles

William V. R. Roberts

Independent Work

Program in Applied and Computational Mathematics

Advised by Professor Mikko Haataja

This paper represents my work in accordance with University regulations.

/s/ William Roberts

1 Introduction

The study of active matter considers non-equilibrium systems of particles where individual particles are self-driven, continuously consuming energy and therefore violating time-reversal symmetry. Some examples of active matter systems include bacteria and other self-propelled organisms or microscale robotic swimmers. The non-equilibrium dynamics of these active systems can result in unique behavior not seen in equilibrium systems, such as motility-induced phase separation (MIPS), in which repulsive and motile (self-propelled) particles spontaneously undergo phase separation and clump up without an attractive term. A commonly used approach to model systems like these is to describe and evolve the system using a continuous density field that varies in space and time rather than tracking innumerable individual particles. [1]

A natural foundation for understanding active dynamics is "Model B", which is a description of the passive phase separation of a conserved scalar field $\phi(r,t)$ with a Landau free energy. The free energy is constructed with a double well: one at a concentrated value for ϕ and one at a dilute value $(-\alpha\frac{\phi^2}{2}+\beta\frac{\phi^4}{4})$.[1] This drives the system to phase separate. The Model B free energy also includes a squared gradient term ($|\nabla\phi|^2$) penalizing sharp spatial derivatives, which discourages interfaces from becoming nonphysically steep. The chemical potential is obtained by taking the functional derivative of the free energy. This potential, therefore, includes a term from the bulk double-well free energy that drives phase separation and a term from the squared gradient term that smooths nonphysical sharpness. The system evolves by pushing density down the gradients of this potential, which conserves material and results in phase separation. [2]

An extension, "Active Model B" (AMB) retains the free energy terms from the passive Model B but adds to the chemical potential a new term proportional to the squared gradient of the local density field ($\lambda|\nabla\phi|^2$). This "lambda" term cannot be written as a functional derivative of any free energy term, meaning that its inclusion is a way to break time reversal symmetry while maintaining mass

[1]The inclusion of only even terms is a natural result of our desire to keep the free energy symmetric. Our choice of -1 or +1 for dense ϕ is arbitrary, so we should not expect inverting that choice to have implications on the free energy.

conservation. The breaking of TRS means AMB is a non-equilibrium, or active, system, although the lambda term does little to actually affect the actual behavior of the system. [2]

In order to see real physical differences, "Active Model B +" (AMB+) extends AMB by adding another TRS-breaking term. The new term $\zeta(\nabla^2\phi)\nabla\phi)$ is added to the current term, meaning we are introducing an interface current term that is not bound to follow the gradient. Before the addition of the ζ term, the long term behavior of Models B and AMB are governed by a process known as Ostwald ripening, which pushes the system towards unlimited phase separation. The interface term $|\nabla\phi|^2$ is stronger for regions of greater curvature, meaning smaller droplets are less energetically favorable than larger droplets. This drives the gradual consolidation of density into large single-phase domains and eliminates smaller ones. However, the ζ term in AMB+ inhibits this behavior, as it causes material to flow away from larger regions to smaller droplets. This is known as reverse Ostwald ripening, and all the terms together result in system that settles to a non-equilibrium steady state droplet size that is sustained by continuous energy consumption.

In summary, we begin with the following AMB+ model:

$$\mathcal{F} = \int \left(-\frac{a}{2}\phi^2 + \frac{b}{4}\phi^4 + \frac{\kappa}{2}|\nabla\phi|^2 \right) d\mathbf{r} \tag{1}$$

$$\mu = -a\phi + b\phi^3 - \kappa\nabla^2\phi + \lambda|\nabla\phi|^2 \tag{2}$$

$$\partial_t\phi = \nabla \cdot \left[M \left(\nabla\mu - \zeta(\nabla^2\phi)\nabla\phi \right) \right] \tag{3}$$

Here M in the current term is the mobility, which characterizes the ease with which matter can flow. In a simplified AMB+ model, M is often treated as a constant scalar for computational ease. We begin with an existing AMB+ time-stepping model constructed by Joshua Arrington. The initial model contains a constant scalar M and no wall interactions. The aim of this paper is to extend

the model to include solid obstacles, including wall free energy and wetting terms, along with a spatially varying tensor mobility.

2 Model and Additions

2.1 Wall Potential

To introduce a solid obstacle into the simulation domain, we first establish the scalar field $w(r)$ once at the beginning of the simulation. The field is designed such that it takes a value of 1 when inside the solid bulk of the wall, a value of zero when outside, and a smooth transition between the two utilizing a hyperbolic tangent function. The smoothness is important to avoid computationally stiff changes and to avoid Gibbs' artifacts when we later move to Fourier space. The wall function is different for every desired shape, but we provide the most basic solid circular shape below as an example:

$$w(\mathbf{r}) = \frac{1}{2}\left(1 + \tanh\left(\frac{-(r-R)}{w_{width}}\right)\right) \tag{4}$$

Here, R is the radius of the circle and w_{width} is a parameter that controls the steepness of the wall. Figure 1 shows some sample cross sections as w_{width} varies. If w_{width} grows too large compared to R, then the hyperbolic tangent never goes properly flat.

The wall field is incorporated into the dynamics through a general wetting free energy, taken from Davis et al.:

$$f_{wet} = -\xi_1\phi|\nabla w|^2 + \xi_2 w\left(\phi - \frac{\phi^3}{3}\right) \tag{5}$$

[3]

Here, ξ_1 and ξ_2 are tunable parameters for their respective terms. The first term is relevant on the interface, where $|\nabla w|^2$ is large, and controls the attraction between the dense phase and the wall. If ξ_1 is positive, the dense phase is attracted while the inverse is true for a negative ξ_1. The second term penalizes ϕ for taking on

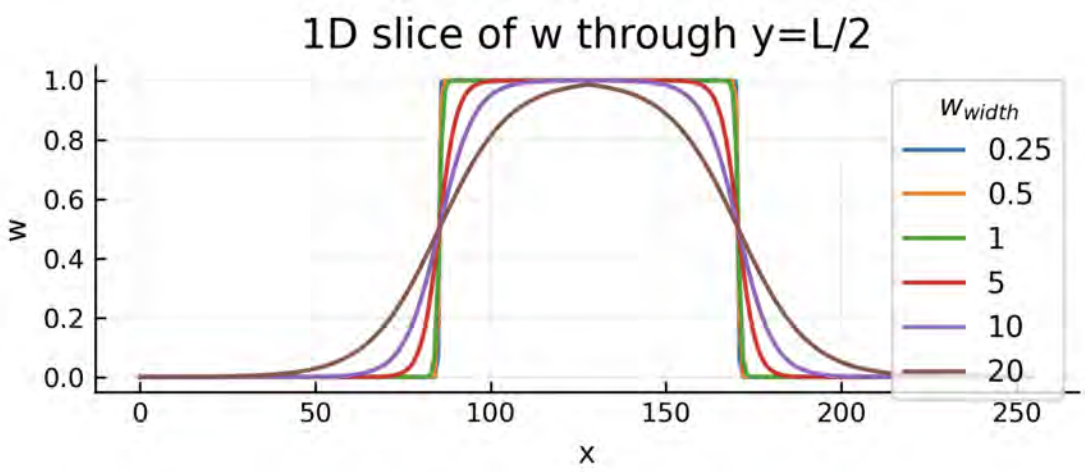

Figure 1: Cross sections of a centered w(r) along the center line of domain while wall width varies.

dense phase values inside the wall when ξ_2 is positive. Together, these parameters and terms create a highly tunable model for wall-density field interaction.

The functional derivative of f_{wet} is then simply added to the existing chemical potential μ.

$$\frac{\delta f_{wet}}{\delta \phi} = -\xi_1 |\nabla w|^2 + \xi_2 w(1-\phi^2) \tag{6}$$

$$\mu = -a\phi + b\phi^3 - \kappa\nabla^2\phi + \lambda|\nabla\phi|^2 + \frac{\delta f_{wet}}{\delta \phi} \tag{7}$$

2.2 Variable Mobility

In a simpler AMB+ model, the mobility M is often treated as a constant scalar, which means material transport at every point and direction in the domain is equally and isotropically easy. This assumption is numerically advantageous, as M can then be brought outside of the ∇ operator in the equation of motion. This results in $\partial_t\phi = M(\nabla^2\mu - \nabla\cdot\zeta(\nabla^2\phi)\nabla\phi)$, which simplifies computation. This assumption is generally accurate in bulk regions, but is nonphysical near a solid obstacle, when the mobility should be neither constant nor isotropic.

In a domain with obstacles, we need M to vanish inside the wall, recover a scalar M away from the wall, and have tangential but not normal mobility along the wall interface. This suppresses mobility in the direction normal to the wall, which is physically accurate.

To create this mobility, we introduce a spatially varying 2 x 2 tensor $\mathbf{M}$(r) at every point in the domain.

$$\mathbf{M}(\mathbf{r}) = M\left(\mathbf{I} - \left(|\nabla w|_{norm}\right)\hat{n} \otimes \hat{n}\right)(1 - w) \tag{8}$$

Here, $|\nabla w|_{norm}$ is the normalized magnitude of the gradient of the wall field $w(r)$, $\hat{n}$ is the wall normal $\frac{\nabla w}{|\nabla w|}$, and the outer product $\hat{n} \otimes \hat{n}$ projects onto the wall-normal direction. The $|\nabla w|_{norm}$ is zero in the bulk away from walls, meaning $\mathbf{M}(r) = \mathrm{M}\,\mathbf{I}$. In the wall, $(1 - w)$ sets $\mathbf{M}$(r) = 0. Finally, along the wall interface, $\mathbf{I}$ - $\hat{n} \otimes \hat{n}$ allows tangential mobility without normal mobility. As a validation of this approach, Figure 2 shows the elements of $\mathbf{M}$ when applied to a circular wall. As we expect, the [0,0] element of $\mathbf{M}$, which covers mobility in the x direction, decreases as we approach the wall from the x direction, but remains constant as we approach from the y direction. The inverse is true for the [1,1] element, which handles y mobility. The rightmost subplot shows the mobility for only one of the off-diagonal terms, but we can see off-diagonal mobility decreasing when normal to the circle and increasing when tangential.

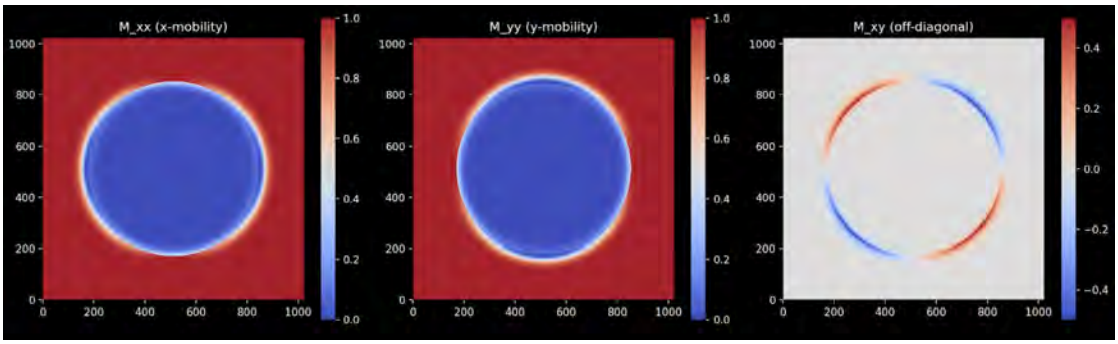

Figure 2: Validation of mobility elements with a circular $w(r)$.

3 Numerical Approaches

3.1 Fourier

When the wall terms and spatially varying mobility are included into our model, the Euler step becomes:

$$\phi^{n+1} = \phi^n + \Delta t \nabla \cdot \mathbf{M} \cdot \left(\nabla \mu^n - \zeta (\nabla^2 \phi^n) \nabla \phi^n \right) , \tag{9}$$

where

$$\mu^n = a\phi^n + b(\phi^n)^3 - \kappa \nabla^2 \phi^n - \xi_1 |\nabla w|^2 + \xi_2 w(1 - (\phi^n)^2) + \lambda |\nabla \phi^n|^2 \tag{10}$$

An issue here is the presence of several higher-order derivative terms, especially a $\kappa \nabla^4 \phi$ term that comes from taking the divergence of the gradient of the already second-order derivative $|\nabla \phi|^2$ term. This 4th order derivative term tricky to numerically compute, and can cause instability or numerical artifacts. A trick known as the spectral method avoids this issue by representing ϕ in Fourier space. In Fourier space, a derivative is simply found by multiplying a term by $i\mathbf{k}$, where $\mathbf{k}$ is the wavenumber set by the grid geometry. Notably, some terms cannot be easily calculated in Fourier space, including the non-linear terms λ and ζ terms. Therefore, those terms are calculated in real space then transformed back into Fourier space for the outside divergence calculation.

We therefore reduce our Euler step in real space to:

$$\hat{\phi}^{n+1} = \hat{\phi}^n + \Delta t \cdot i\mathbf{k} \cdot \mathcal{F} \left\{ \mathbf{M} \cdot \left(\mathcal{F}^{-1} \left[i\mathbf{k} \cdot \hat{\mu}^n \right] - \zeta (\nabla^2 \phi^n) \nabla \phi^n \right) \right\} , \tag{11}$$

where

$$\hat{\mu}^n_{soft} = \mathcal{F} \left[a\phi^n + b(\phi^n)^3 - \xi_1 |\nabla w|^2 + \xi_2 w(1 - (\phi^n)^2) + \lambda |\nabla \phi^n|^2 \right] + \kappa k^2 \hat{\phi}^n \tag{12}$$

3.2 Semi-Implicit

While the spectral method reduces the computational demands of calculating higher-order derivatives, it introduces a severe limitation in maximum allowable time step Δt. The κk^4 term grows very quickly for the largest wavenumbers, which causes instabilities in the Euler steps. [4]

Using the same stability condition as in Zhu et al. (1999), we can set a limit at:

$$\Delta t < \frac{1}{\kappa k_{max}^4} \tag{13}$$

This limits Δt to such a small value that reaching long enough simulation times to see phase separation or coarsening effects can be impractical. To avoid this situation, we use a semi-implicit scheme as introduced in Zhu et al. (1999) [4]. Starting from the explicit Fourier update equation above, we first add $(\alpha \Delta t \kappa k^4)\hat{\phi}^{n+1}$ to the left-hand side and $(\alpha \Delta t \kappa k^4)\hat{\phi}^{n}$ to the right-hand side, which is valid when Δt is small enough that $\hat{\phi}^n \approx \hat{\phi}^{n+1}$. Here, α is a tunable parameter that governs the stabilization of the semi-implicit method. We follow Zhu et al. in setting $\alpha = 0.5$, but it can be adjusted to increase or decrease the damping of the semi-implicit scheme. [4]

Finally, both sides of the update can be divided by $(1 + \alpha \Delta t \kappa k^4)$, which recovers $\hat{\phi}^{n+1}$ on the left, giving us the final semi-implicit Fourier update:

$$\hat{\phi}^{n+1} = \hat{\phi}^n + \frac{\Delta t \cdot i\mathbf{k} \cdot \mathcal{F}\left\{\mathbf{M} \cdot \left(\mathcal{F}^{-1}\left[i\mathbf{k} \cdot \hat{\mu}^n\right] - \zeta(\nabla^2 \phi^n)\nabla\phi^n\right)\right\}}{1 + \alpha \Delta t \kappa k^4} \tag{14}$$

Crucially, we now have a $\Delta t \kappa k^4$ term in both the numerator and denominator, canceling out the effects of high-k growth, and allowing us to use much larger values for Δt.

4 Results

4.1 Recovery of Model B

Our first attempt at validating the new approach is to recover the behavior of the passive Model B when we set the active terms λ and ζ to zero. Because the wetting terms act only along or within the wall interface, their value does not affect this check. As explained in the introduction, the passive model should result in a gradual coarsening of the phases driven by the free energy.

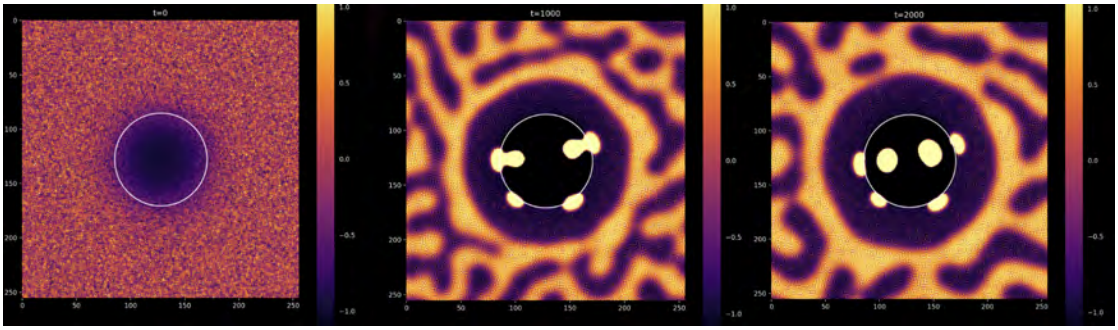

Figure 3: Recovery of passive Model B Coarsening. Panels show t = [0, 1000, 2000] $\lambda = 0$, $\zeta = 0$, $\xi_{1,2} = 4.5$

As we can see in Figure 3, setting the active parameters to 0 recovers the expected passive Model B behavior. The initial random distribution of ϕ separates into dense and dilute phases, which coarsen over time as predicted by Ostwald ripening. Note the wetting parameters $\xi_{1,2}$ were set at a very high value of 4.5, which creates a very dilute phase around the wall with several extremely dense droplets on the interface. Given enough time, these droplets migrated into the interior of the wall, a drawback to the hyperbolic tangent nature of the wall values.

We can prove the recovery of passive behavior more rigorously by tracking the domain scale L over time, as is done in Cates and Tailleur's MIPS Review. [1] They show an L $\sim t^{1/3}$ relation at longer time scales for particles in AMB. While AMB is active, the absence of the ζ term from AMB+ means the behavior looks like a passive Model B. Therefore, for this passive test, we should see the continued coarsening in the phases following a $\sim t^{1/3}$ relation at high time scales.

In these plots of domain size over time, we see the expected behavior once t

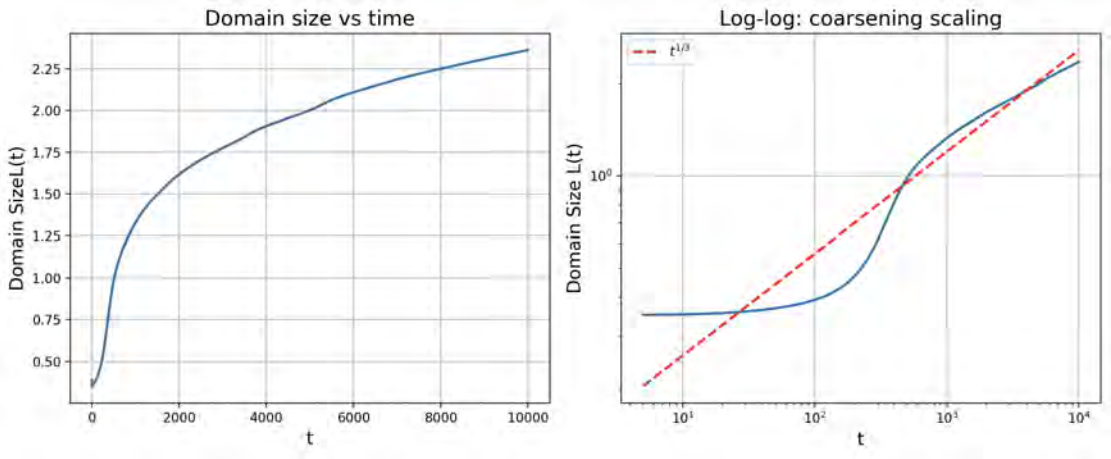

Figure 4: Domain Scale L(t) over time for $\lambda = 0$, $\zeta = 0$, $\xi_{1,2} = 4.5$

exceeds about 200, as predicted in the MIPS review. As a note, we calculated $L(t)$ by finding the mean wavenumber in the bulk areas. The inverse of this wavenumber gives us L.[2] Looking carefully at the high time values of the right subplot, it is possible that the slope may be falling below the expected power scaling. It may be worth running the Euler steps for even longer in order to characterize this behavior or rule it out. It may also be possible that my calculation for L is distorted by the presence of the circular wall.

4.2 Active Effects

We are similarly interested in what the behavior looks like in an active context with nonzero λ and ζ terms. As we can see in Figure 5 when λ and ζ were both set to 0.8, the coarsening behavior seems to stop once it reaches a certain level. The central and right subplots show little additional coarsening, suggesting the influence of reverse Ostwald ripening.

We can, as with the passive case, quantify this coarsening with an examination of $L(t)$.

As we can see in Figure 6, the originally rapid growth in L at low t quickly flattens out as the log-log slope falls away from 1/3. Therefore, we can confirm

[2]The design of this wavenumber selection was suggested by generative AI.

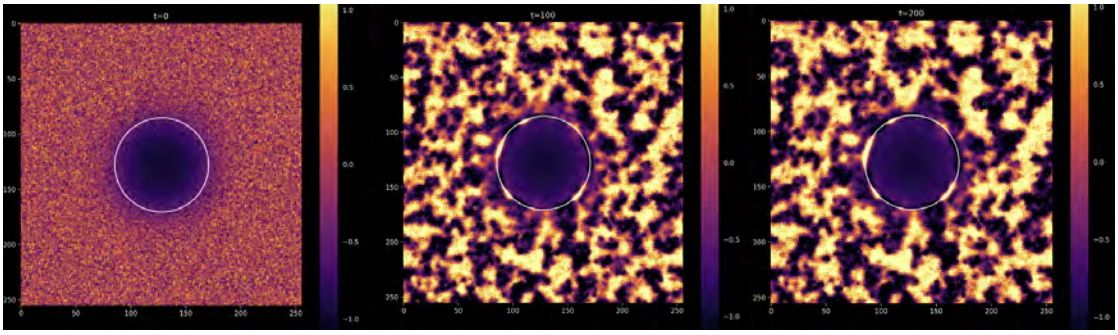

Figure 5: Active AMB+ over time for $\lambda = -0.8$, $\zeta = 0.8$, $\xi_{1,2} = 0.5$

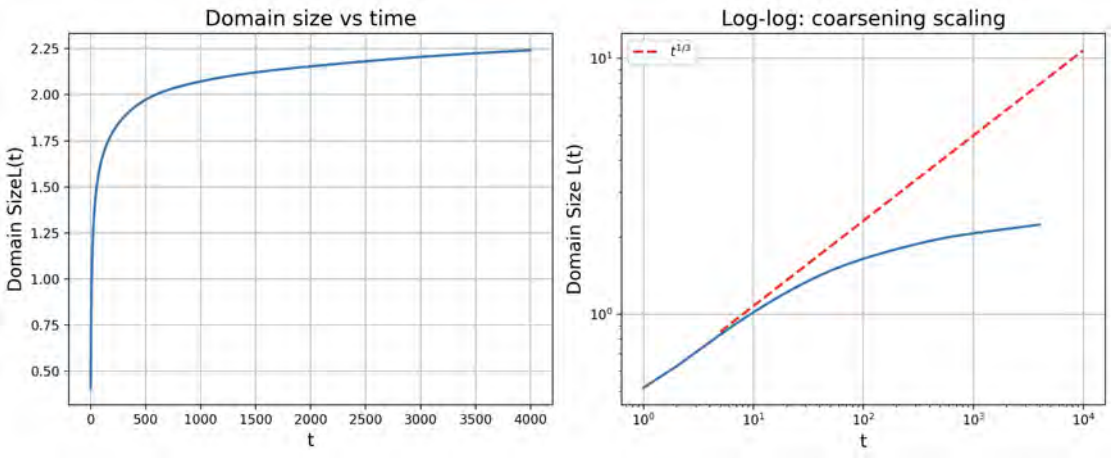

Figure 6: Domain Scale L(t) over time for $\lambda = 0.8$, $\zeta = 0.8$, $\xi_{1,2} = 0.5$

the correct presence of active terms driving reverse Ostwald ripening, and judge the bulk effects of this model to be a success. It is interesting to note, though, that at low time-scales the growth in domain size is significantly more rapid than the passive case. This suggests the ζ term may actually drive the consolidation of very small regions, before its contribution to reverse Ostwald ripening begins to take effect at larger sizes.

subsectionWetting We are also interested in the efficacy of the wetting terms implemented as the wall potential. Unfortunately, in an active system with non-zero λ and ζ terms, we found that the magnitude of the active ζ term seemed to overwhelm wetting effects, meaning varying the parameters did little to change the wetting behavior. However, in a passive system, we were able to observe the

impacts of the wetting parameters.

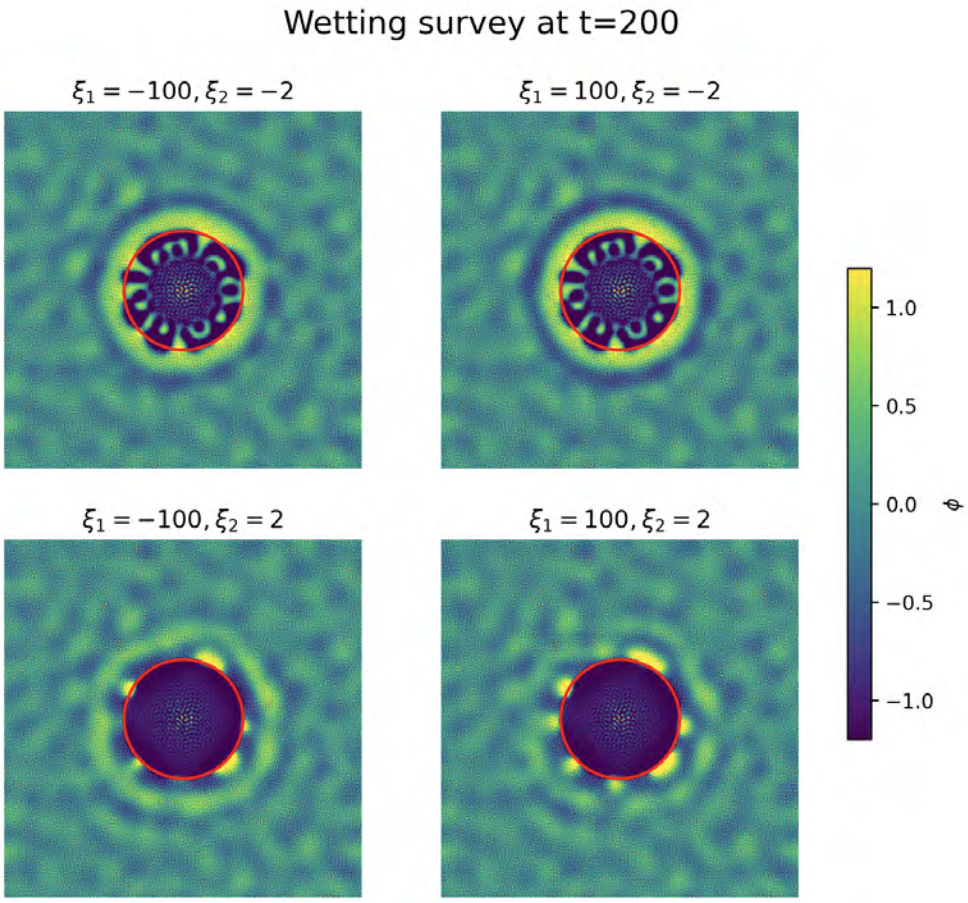

Figure 7: Wetting behavior in passive system as $\xi_{1,2}$ vary. λ, $\zeta = 0$

Figure 7 shows the fairly short term wetting behavior for a set of 4 pairs of ξ_1, ξ_2 values. Attractive wetting seemed to be strongest when ξ_2 was negative, while a positive ξ_2 attracted small pockets of density without the whole ring we see in the top row. Varying ξ_1 had less of an impact, even when its values were dramatically increased, but we can see a faint difference between the bottom left and right subplots, where a negative ξ_1 appears to draw in a faint dense ring. This pattern for ξ_1 contradicts the assumed behavior when establishing the wetting terms, so deserves further investigation. We suspect the weak dependence on ξ_1 occurs because the wall width must be kept fairly wide to minimize Gibbs' artifacts in the center of the wall due to its steepness. Therefore, ∇w never becomes very

large and weakens the ξ_1 term.

5 Conclusion

In this paper, we have derived a new numerical model for AMB+ dynamics when we cannot consider M a constant scalar because of the presence of obstacles. We have shown through recovery of passive models as well as a qualitative comparison to the predicted behavior that this method is accurate, at least in bulk regions. We have also implemented a wetting scheme that can successfully attract dense phases (specifically when ξ_2 is negative) in passive systems, and shows some signs of droplet-like wetting at positive values of both ξ parameters. More work should be done to extend this wetting characterization to the system when λ, ζ are not equal to zero and to find relevant values for the tunable parameters ξ_1, ξ_2 in an active regime.

References

[1] Michael E. Cates and Julien Tailleur. Motility-induced phase separation. *Annual Review of Condensed Matter Physics*, 6:219–244, 2015.

[2] Elsen Tjhung, Cesare Nardini, and Michael E. Cates. Cluster phases and bubbly phase separation in active fluids: Reversal of the ostwald process. *Physical Review X*, 8, 2018.

[3] Ryan Davis, Fadi Abdeljawad, Jeffrey Lillibridge, and Mikko Haataja. Phase wettability and microstructural evolution in solid oxide fuel cell anode materials. *Acta Materialia*, 78:271–281, 2014.

[4] Jingzhi Zhu, Long-Qing Chen, Jie Shen, and Veena Tikare. Coarsening kinetics from a variable-mobility Cahn-Hilliard equation: Application of a semi-implicit Fourier spectral method. *Physical Review E*, 60:3564–3572, 1999.

Acknowledgments

Thank you to Professor Haataja and Josh Arrington for their wonderful help on this interesting project.

Generative AI was used for Latex formatting, figure generation, and for a method to find the dominant wavenumber for $L(t)$