

A Markov Chain Model of the Viral Spread of Internet Memes

Rebecca Robinson

April 18, 2017

Student Research Conference, 2017
University of Michigan-Flint

Contents

1	Introduction	3
2	Setting up the Markov Chain Model	3
2.1	Defining Our States	3
2.2	Describing Our Probabilities and Forming Our Markov Chain	3
3	Getting Our $I(t)$ Equation	4
3.1	Diagonalizing M	4
3.2	Our Equations	5
4	Regression Analysis	5
5	Can Calculus Tell Us Anything?	7
6	Predictions	7
6.1	Predicting α	8
6.2	Predicting β	9
6.3	Predicting γ	10
7	Appendix	11
7.1	Examples Used	11
7.1.1	Dat Boi	11
7.1.2	Doge	12
7.1.3	“Flint, Michigan”	12
7.1.4	Damn Daniel	13
7.1.5	Nyan Cat	14
7.1.6	Trash Dove	14
7.1.7	Impossibru	15
7.1.8	“Rick Roll”	16
7.1.9	Harlem Shake	16
7.1.10	“ermahgerd”	17
7.2	Parameter Regression	18

1 Introduction

With the rising use and popularity of the Internet, it is much easier to spread content from person to person with speed. Content in the form of humorous images known as **memes** are often shared among peers. In this project, we aim to model the rise and fall of Internet memes by using methods normally used to model infectious diseases. We first provide a Markov chain version of an SIR model, and then use that to describe a meme's popularity over time, with the ultimate goal of wanting to predict the popularity given an initial set of data points.

2 Setting up the Markov Chain Model

2.1 Defining Our States

The inspiration for this project comes from an paper where the author used SIR models and differential equations to model the spread of a meme. In this paper, we attempt to do the same using Markov chains instead of differential equations.

In setting up the Markov chain model, we imagine the spread of a meme from person to person as the spread of a disease. Let's carefully describe our states. We have three states in our Markov chain: **susceptible**, **infected**, and **recovered**, from which the term SIR models comes from. A person in the susceptible state is one that has not been exposed to the meme and is unaware of it. A person in the infected state is a person that has seen the meme and decides to actively share the meme with peers, and a person in the recovered state is one who is no longer sharing the meme and is no longer interested in it.


From the susceptible state, an individual can only move to the infected state. In the event, that an individual sees the meme and decides not to share it, we consider that to be essentially the same as if they did not see the meme at all, so they remain in the susceptible state. From the infected state, an individual can only move to the recovered state. That is, an individual sharing the meme can only become disinterested in the meme (they cannot "unsee" the meme and move back to the susceptible state). Similarly, an individual in the recovered state also cannot move back to the susceptible state, but they can move back to the infected state. In other words, this means that they become reinterested in the meme.

2.2 Describing Our Probabilities and Forming Our Markov Chain

We define three different variables that we will be using in our Markov chain.

1. α : The proportion of susceptible individuals that become infected after one time step.
2. β : The proportion of infected individuals that become recovered after one time step.
3. γ : The proportion of recovered individuals who become re-interested in the meme and become infected again.

We create the following Markov chain diagram, with probabilities labeled. Here, S, I, and R represent the susceptible, infected, and recovered states respectively.



With this Markov chain, we create the following transition matrix.

$$M = \begin{bmatrix} 1-\alpha & 0 & 0 \\ \alpha & 1-\beta & \gamma \\ 0 & \beta & 1-\gamma \end{bmatrix}$$

Now that we have our transition matrix M , we can start with the fun stuff.

3 Getting Our $I(t)$ Equation

3.1 Diagonalizing M

For this project, what we're really interested in is the probability of being in the infected state after n time steps. In order to find the proportion of individuals in each state after an certain time n , we want to find M^n . Each $M_{i,j}$ entry in this matrix will give us the probability (or proportion of individuals) that end up in state i from state j .

First, we find the eigenvalues. Our eigenvalues for this matrix are: $1 - \alpha$, $1 - \beta - \gamma$, and 1 . Then we find the corresponding eigenvectors for each eigenvalue:

$$\text{For } 1 - \alpha: \begin{bmatrix} 1 \\ \frac{\gamma - \alpha}{\alpha - \beta - \gamma} \\ \frac{\beta}{\alpha - \beta - \gamma} \end{bmatrix} \quad \text{For } 1: \begin{bmatrix} 0 \\ 1 \\ \frac{\beta}{\gamma} \end{bmatrix} \quad \text{For } 1 - \beta - \gamma: \begin{bmatrix} 0 \\ 1 \\ -1 \end{bmatrix}$$

We combine these together to make a new matrix, which we'll call P , and then we also create the matrix of our eigenvalues, D .

$$P = \begin{bmatrix} 1 & 0 & 0 \\ \frac{\gamma - \alpha}{\alpha - \beta - \gamma} & 1 & 1 \\ \frac{\beta}{\alpha - \beta - \gamma} & \frac{\beta}{\gamma} & -1 \end{bmatrix}$$

$$D = \begin{bmatrix} 1 - \alpha & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 - \beta - \gamma \end{bmatrix}$$

We leave it to the reader to check that $M = PDP^{-1}$. Because of this, now we have that $M^n = PD^nP^{-1}$. The value that we are the most concerned with is $M_{2,1}^n$, which is the proportion of the population that is "infected" with the meme at the n^{th} time step. Because we now have a way to calculate this value with our diagonalization, let's find an explicit formula for the proportion of infected individuals depending on n, α, β , and γ . Since we want to think of this function as a continuous function, we'll represent time as t , rather than the discrete time step n .

3.2 Our Equations

By multiplying these matrices out, we can find explicit formulas for the proportion of susceptible individuals, the proportion of infected individuals, and the proportion of recovered individuals. We present them below:

$$S(t) = (1 - \alpha)^t$$

$$I(t) = \frac{\alpha\beta(1 - \beta - \gamma)^t + \gamma(\alpha - \beta) - \gamma^2 - (\alpha\beta + \gamma(\alpha - \beta) - \gamma^2)(1 - \alpha)^t}{\alpha\beta - \beta^2 + \gamma(\alpha - 2\beta) - \gamma^2}$$

$$R(t) = \frac{\beta(1 - \alpha)^t}{\alpha - \beta - \gamma} + \frac{\beta}{\beta + \gamma} - \frac{\alpha\beta(1 - \alpha - \beta)^t}{\alpha\beta - \beta^2 + \gamma(\alpha - 2\beta) - \gamma^2}$$

While it's nice to have $S(t)$ and $R(t)$, what we're really concerned with is $I(t)$. $I(t)$ is the curve that models the popularity of the meme over time. Depending on different α, β , and γ values, we get curves with a variety of different peaks and widths. Through experimentation of changing α, β , and γ , we can see the effect that each variable has on the curve.

First, α controls how quickly the function reaches the maximum. If we increase α , the peak of the curve is higher, and it also occurs at a smaller t . If we think about this, it makes sense: α is the probability that a person moves from the susceptible state to the infected state. A higher α would mean more people move from the susceptible state to the infected state at each time step.

Now let's look at β . Keeping α and γ constant, changing β changes the width of the peak. A higher β value makes the peak narrower. Does this make sense? If we think about what β is, yes! Recall that β is the probability of moving from infected to recovered. If this is higher, a person is less likely to remain in the infected state from one time step to the next. Because of this, the proportion of infected individuals decreases quicker the higher β is.

Last but not least, let's consider γ . When we change γ , we change the asymptote that our $I(t)$ has. Raising the γ value lifts the asymptote up. If we think about what γ is, this makes perfect sense. Since γ is the probability that an individual will become re-interested in the meme, raising it would increase the proportion of individuals that remain at the infected state over time.

We're also interested in the long term behavior of our $I(t)$ function. From the steady-state distribution of our Markov chain, we can determine that the $I(t)$ function has a horizontal asymptote at $\frac{\gamma}{\beta + \gamma}$. We can verify this by looking at the limit as $t \rightarrow \infty$.


$$\lim_{t \rightarrow \infty} I(t) = \frac{\gamma(\alpha - \beta - \gamma)}{(\beta + \gamma)(\alpha - \beta - \gamma)} = \frac{\gamma}{\beta + \gamma}$$

This is useful as it tells us our long term behavior for our $I(t)$, and tells us where the popularity ultimately stabilizes at. Given our β and γ we can find exactly this.

4 Regression Analysis


We want to find the α, β, γ that change the model to best fit our given data. We obtain our data from Google Trends.

First, we assume that the popularity of a meme is directly related to the popularity of it as a search term. With this assumption, we can take data from Google Trends, which gathers information about the popularity of search terms. Let's take a look of Google Trends curve for the popular meme "dat boi".



Before we go into detail about getting $I(t)$ from this data, we should note a couple things about the Google Trends. The y -axis is scaled so that 100 denotes the peak popularity of the search term, with the rest of the data being scaled from there. Using SageMath, we can perform a regression analysis on points from the Google Trends data, and find our α, β, γ values that best fit our $I(t)$ equation to the data. First, though, we must transform our data in a way that makes sense. The values that Google provides us are in terms of "popularity index", with a new data point every week. Because of this, we make our time step equal to one week. Additionally, we must transform that "popularity index" to the proportion of the population that is infected. We make the base argument that we'll say a meme has hit peak popularity when it reaches 60% of the population, and we'll scale the rest of the data to match this.

Plotted below is data points from the "dat boi" Google Trend graph, along with our best fit curve from our regression analysis.



This best fit line corresponds to $\alpha \approx 0.5475$, $\beta \approx 0.3630$, and $\gamma \approx 0.1063$. This tells us a little bit more about the spread of the meme through the population. Recalling what α, β, γ are, every

susceptible individual has a 54.75% chance of becoming “infected” with the meme, every infected individual has a 36.3% chance of becoming disinterested in the meme, and every recovered individuals 10.63% of becoming reinterested in the meme after each time step. In the case of this particular meme, our time steps are a week in length.

5 Can Calculus Tell Us Anything?

We’ve now found an equation $I(t)$ that models the spread of a meme through a population. Using calculus, we can learn more information about $I(t)$. We start by differentiating with respect to t .

$$I'(t) = \frac{\alpha\beta \log(1-\beta-\gamma)(1-\beta-\gamma)^t - (\alpha\beta + \gamma(\alpha-\beta) - \gamma^2) \log(1-\alpha)(1-\alpha)^t}{\alpha\beta - \beta^2 + \gamma(\alpha + 2\beta) - \gamma^2}$$

Then, with this, we can find an explicit equation the time we have our maximum in terms of α, β , and γ .

$$t = \frac{\log\left(\frac{\alpha\beta \log(1-\beta-\gamma)}{(\alpha\beta + \gamma(\alpha-\beta) - \gamma^2) \log(1-\alpha)}\right)}{\log\left(\frac{1-\alpha}{1-\beta-\gamma}\right)}$$

Additionally, we can also find our second derivative.

$$I''(t) = \frac{\alpha\beta \log^2(1-\beta-\gamma)(1-\beta-\gamma)^t - (\alpha\beta + \gamma(\alpha-\beta) - \gamma^2) \log^2(1-\alpha)(1-\alpha)^t}{\alpha\beta - \beta^2 + \gamma(\alpha + 2\beta) - \gamma^2}$$

We can also find our time for the inflection point of $I(t)$, which would be when it is decreasing the fastest.

$$t = \frac{\log\left(\frac{\alpha\beta \log^2(1-\beta-\gamma)}{(\alpha\beta + \gamma(\alpha-\beta) - \gamma^2) \log^2(1-\alpha)}\right)}{\log\left(\frac{1-\alpha}{1-\beta-\gamma}\right)}$$

With our model, we can use our first and second derivatives to find the time of the peak popularity and to find the time where the popularity is decreasing the fastest. Let’s use these with our model of the “dat boi” meme.

With our values for α, β , and γ determined from the regression, we obtain a maximum popularity at $t = 1.779$. Because our time step is a week, this means it reaches its peak popularity during the second week. Additionally, it decreases the fastest at $t = 3.457$.

What does this tell us? Using calculus, we can now find *exactly* where the maximum will occur, instead of judging from the data points where the popularity for a whole week is summarized in one data point. Additionally, we can look at where the function is decreasing the fastest, we can see exactly where the meme’s popularity falls the fastest as well.

6 Predictions

We’ve done some calculus work and on the surface it seems like the calculus might be a good way to start thinking of some ways to predict what the $I(t)$ might be given data points. However, it’s hard to predict what $I(t)$ might be given data points using calculus because we can’t take the derivative or second derivative when we don’t initially know what the function is. So we must look at other ways to predict α, β , and γ .

As a summary, and since we'll refer to it later, let's make a table of different memes/search terms and their respective α, β, γ values. For these, all time steps were weekly data with exceptions. Those marked with an asterisk (*) were computed with daily time steps, and those marked with a dagger (†) were computed with monthly time steps. Trend and regression graphs are provided in Section 7.1.

Meme/Search Term	α	β	γ
Dat Boi	0.535	0.3308	0.0385
Doge†	0.1148	0.0656	0.0034
"Flint, Michigan"*	0.1884	0.1547	0.01392
Damn Daniel*	0.3117	0.22	0.0022
Nyan Cat†	0.3919	0.1842	0.400
Trash Dove*	0.3428	0.2826	0.0011
Impossibru†	0.3857	0.2614	0.0602
"Rick Roll"†	0.4175	0.5357	0.1072
Harlem Shake	0.4039	0.3050	0.0089
"ermahgerd"	0.3239	0.2832	0.0098

The graphs of these different lines plotted against their respective data are in SEction 7.1.

6.1 Predicting α

The probability α is probably the easiest of the probabilities to estimate from initial data points, as its main effect occurs during the beginning of a meme's spread. We can approximate α by taking the maximum of the slopes between two consecutive points. This makes sense when we consider what exactly α is. Since α dictates how quickly the meme will jump in probability, looking at the slopes between the data points would be the best way to guess what α is. However, with this method of approximation, we need to assume that the

We present some examples below that show the estimated α compared to the actual α obtained from the regression.

Meme or Search Term	Estimated α	Actual α	Difference
"dat boi"	0.516	0.535	0.019
Doge	0.174	0.1148	0.0592
"Flint, Michigan"	0.18	0.188	0.008
"Damn Daniel"	0.258	0.3117	0.0537
Nyan cat	0.246	0.3919	0.1459
Trash Dove	0.246	0.3428	0.0968
Impossibru	0.33	0.3857	0.0557
"Rick Roll"	0.24	0.4175	0.1775
Harlem Shake	0.306	0.4039	0.0979
"ermahgerd"	0.222	0.3239	0.1019

Initially, this seemed promising; our first four memes on the list have actual α 's that are close to our estimated α 's, especially in the "Flint, Michigan" case. However, we can see that it is not a good estimate for other memes, notably Nyan Cat and "Rick Roll". Our α prediction method is

almost always an underestimate. This means that α must be higher than the largest slope between two consecutive data points. We could potentially try in the future to see if there is a correlation between the highest slope of two consecutive points and the actual α value and see if we can predict α from this slope value.

So we perform regression analyses by SPSS on the estimated α and the actual α values putting the estimated α as the independent variable and the actual α value as the dependent variable. We present the results below:

Model Summary and Parameter Estimates									
Dependent Variable: actual									
Model Summary						Parameter Estimates			
Equation	R Square	F	df1	df2	Sig.	Constant	b1	b2	b3
Linear	.660	15.542	1	8	.004	.074	.984		
Logarithmic	.767	26.376	1	8	.001	.791	.332		
Quadratic	.797	13.729	2	7	.004	-.351	3.775	-4.037	
Cubic	.797	13.729	2	7	.004	-.351	3.775	-4.037	.000
Exponential	.490	7.701	1	8	.024	.133	3.185		

The independent variable is estimated.

We will take the quadratic regression to be the best regression method, since its R Square value is the highest. We also check that it shows a significant relationship between α_{est} (the estimated α through the slopes) and α . Through the F-Test of Overall Significance, we have a significance level of 0.004, so $p < 0.05$ and we can conclude we have a significant relationship. The R Square value of 0.797 tells us that 79.7% of the variance in the α data is explained by the α_{est} data. Thus, we can present α by the following equation:

$$\alpha = -4.037\alpha_{est}^2 + 3.775\alpha_{est} - 0.351$$

This provides us with a better approximation of α than just the highest slope between consecutive points.

6.2 Predicting β

We notice a positive correlation between α and β , in that with a larger α value, we have a higher β value. Using SPSS, we test our examples and perform various regression analyses on the α and β pairs, with α as the independent variable.

We present the following output from SPSS after running five different regression analyses.

Model Summary and Parameter Estimates

Dependent Variable: beta

Equation	R Square	Model Summary				Statistics			
		F	df1	df2	Sig.	Constant	b1	b2	b3
Linear	.522	8.741	1	8	.018	.079	.599		
Logarithmic	.376	4.819	1	8	.059	.396	.097		
Quadratic	.523	3.841	2	7	.075	.089	.501	.179	
Cubic	.705	4.777	3	6	.050	.290	-3.309	15.608	-17.242
Exponential	.571	10.637	1	8	.012	.099	2.803		

The independent variable is alpha.

We note that the cubic regression has the highest R Square value, but we must also show that it depicts a significant relationship between α and β . We perform a F-Test for Overall Significance and determine that the significance value (which SPSS has rounded) is 0.0498. Since this is less than 0.05, we can conclude that the cubic regression shows a significant relationship between α and β .

Then this means that 70.5% of the variation in the β data can be explained by α . We can also predict the β values accurately 70.5% of the time with the α values and the following equation:

$$\beta = -17.242\alpha^3 + 15.608\alpha^2 - 3.309\alpha + 0.290$$

We'll compare predicted β values to actual β values below, using the actual α values to plug into the equation.

Meme/Search Term	Expected β	Actual β	Difference
Dat Boi	0.3468	0.3308	0.0160
Doge	0.0897	0.0656	0.0241
"Flint, Michigan"	0.1053	0.1547	0.0494
Damn Daniel	0.2529	0.2200	0.0329
Nyan Cat	0.3526	0.1842	0.1684
Trash Dove	0.2952	0.2826	0.0126
Impossibru	0.3463	0.2614	0.0846
"Rick Roll"	0.3743	0.5357	0.1614
Harlem Shake	0.3636	0.3050	0.0586
"ermahgerd"	0.2698	0.2832	0.0134

This gives us an average difference of 0.06214 from the actual value of β . It's fair to say that this is a fairly accurate approximation of β as long as we are given an accurate α value.

6.3 Predicting γ

Predicting γ appears to be the most difficult of the three. If we are only given an initial set of data points, it is hard to predict γ because γ is involved mostly in the long term behavior. Unfortunately, performing the same regressions as above with both α and β as the independent variables separately does not yield any regression line that gives us a significant relationship. This means that γ is not related to α or β in any significant way. This makes it harder to predict γ given α or β .


Unfortunately, with the sake of time for this project, we do not have enough time to go forward with more attempts to predict γ , but instead look forward to working more with this in the future.

7 Appendix

7.1 Examples Used



7.1.1 Dat Boi

We first provide the Google Trends graph, and then we provide the best fit line plotted against the data from the regression analysis with our model, with α, β, γ listed. We follow this same order for the rest of the examples as well.



The α, β, γ values are as follows: $\alpha = 0.5354, \beta = 0.3308, \gamma = 0.0325$.

7.1.2 Doge




The α, β, γ values are as follows: $\alpha = 0.1148, \beta = 0.0656, \gamma = 0.0034$

7.1.3 “Flint, Michigan”

Interest over time ?







The α, β, γ values are as follows: $\alpha = 0.1884, \beta = 0.1547, \gamma = 0.01392$

7.1.4 Damn Daniel



Interest over time ?



The α, β, γ values are as follows: $\alpha = 0.3117, \beta = 0.2200, \gamma = 0.0022$

7.1.5 Nyan Cat


Interest over time ?




The α, β, γ values are as follows: $\alpha = 0.3919, \beta = 0.1842, \gamma = 0.4000$

7.1.6 Trash Dove

Interest over time ?







The α, β, γ values are as follows: $\alpha = 0.3428, \beta = 0.2826, \gamma = 0.0011$

7.1.7 Impossibru



Interest over time ?



The α, β, γ values are as follows: $\alpha = 0.3857, \beta = 0.2614, \gamma = 0.0602$

7.1.8 “Rick Roll”

Interest over time ?




The α, β, γ values are as follows: $\alpha = 0.4175, \beta = 0.5357, \gamma = 0.1072$

7.1.9 Harlem Shake

Interest over time ?







The α, β, γ values are as follows: $\alpha = 0.4039, \beta = 0.3050, \gamma = 0.0089$

7.1.10 “ermahgerd”


Interest over time ?



The α, β, γ values are as follows: $\alpha = 0.3239, \beta = 0.2832, \gamma = 0.0098$

7.2 Parameter Regression

Below is a graph of the regression performed on the α_{est} and α data. α_{est} is on the x -axis and α is on the y -axis.



Below is a graph of the different regression analyses performed on the α, β pairs with β as the dependent variable and α as the independent variable.

