

A Simple Method for Evolving Large Character Social Networks

James Ryan, Michael Mateas, and Noah Wardrip-Fruin

Expressive Intelligence Studio
University of California, Santa Cruz
{jor, michaelm, nwf}@soe.ucsc.edu

ABSTRACT

We present a method for generating social networks for gameworlds with very many characters. The method, a generalization of the approach we employ in *Talk of the Town*, operates from a simple principle: characters' affinities for one another evolve as a function of the compatibility of their personalities and the amount of time they spend together. By this principle, friendships emerge as compatible characters interact more extensively and enmities emerge as *incompatible* characters do so. Beyond platonic affinity, our method evolves romantic feelings by the same principle (applied to romantic considerations). How exactly compatibility is defined (*e.g.*, by operationalizing psychological, sociological, or artistic theories) may vary according to the goals of a particular project. In this paper, we describe our method in generalized terms that are agnostic to our specific application of it, so that interested readers may implement it in their own systems.

Keywords

world generation, social simulation, affinity modeling, non-player characters, believability

INTRODUCTION

Human beings embed themselves in large, dynamic social networks, but this has not been the case with non-player characters (NPCs) in games. Partly, this is because games often do not feature enough characters across which large social networks may be constituted, but even games with enough characters typically do not feature dynamic social networks. To be clear, we consider a character social network to be *large* if it comprises several dozen or more NPCs, and *dynamic* if there is modeling of the mutable social relationships among those NPCs. Additionally, character social networks should be *coherent*, meaning that the relationships characters have should be believable with respect to the larger gameworld.

Believability is hampered when NPCs in games do not appear to be socially situated (Afonso and Prada, 2008), but large, dynamic, coherent social networks are difficult to hand-author or generate. In this paper, we present a generalized method for generating such networks by modeling them as simple dynamical systems that evolve over time from the interaction of basic social mechanics. This produces rich social networks with limited computational burden. Specifically, our method operates from a straightforward principle: characters' (nonreciprocal) affinities for one another evolve as a function of the compatibility of their personalities and the amount of time that they spend together. From this, friendships emerge as compatible characters interact more extensively and enmities emerge as incompatible characters do so. Beyond platonic affinity, our method evolves romantic feelings by the same principle (applied to romantic considerations). How exactly personality and romantic compatibility is

Proceedings of 1st International Joint Conference of DiGRA and FDG

©2016 Authors. Personal and educational classroom use of this paper is allowed, commercial use requires specific permission from the author.

defined (*e.g.*, by operationalizing psychological, sociological, or artistic theories) may vary according to the goals of a particular project; as we explain below, in our implementation we operationalize findings from the social sciences (Verbrugge, 1977; Luo and Zhang, 2009; Selfhout et al., 2010). We should also note that we do not claim that this model of human affinity is the only viable one—we can also imagine a model in which personality compatibility is not as important as other factors (*e.g.*, important shared experiences, networks of kinship), though such models would perhaps not lend to simple evolution procedures. While we have implemented this method as part of a larger framework, herein we describe it in generalized terms so that interested readers may use it in their own applications.

DESIGN GOAL

The design goal of this project is to develop a generalized, efficient method for producing a fully connected network of affiliations between NPCs, and a history of the development of those affiliations, that can be called upon during gameplay. (In *Talk of the Town*, such networks serve as the structural scaffold for character knowledge propagation, which is our game’s core focus.) While this could be done by handcrafting or by high-fidelity simulation, these approaches would make the iteration that is necessary for game design (and, moreover, the generation of unique social networks for each gameplay instance) infeasible or too computationally expensive. As such, our method utilizes low-fidelity simulation and a simplified notion of character compatibility that is amenable to the operationalization of varying psychological, sociological, and artistic theories. By sacrificing nuance, it becomes feasible to quickly generate large character social networks; further, as we discuss below, emergent dynamics (*e.g.*, from feedback loops) of the simple procedure *do* still yield some richness. Finally, our method can easily be extended to support particular kinds of complexity that are important to a given project; later, we provide examples of extensions that we have included in our own implementation.

RELATED WORK

Several earlier projects, such as *Prom Week* and *Versu*, have modeled mutual character affinities, but across networks of few characters (McCoy et al., 2014; Evans and Short, 2014). While in such systems it has been feasible for authors to specify consistent social networks, in this paper we present a method for generating social networks across hundreds or potentially thousands of characters, a scale at which manual specification is infeasible.

Our method is housed in the larger *world generation* procedure of *Talk of the Town*, outlined in Ryan et al. (2015), which simulates a socially oriented American small town from 1839 until 1979. As such, our approach evokes roguelike techniques (Garda, 2013), and we note that we have specifically been inspired by *Dwarf Fortress* (Adams, 2015), whose approach to world generation is famous. But while world generation in *Dwarf Fortress* integrates myriad concerns—terrain, climate, demography, war, more—that may not be at play in other applications (*e.g.*, games outside fantasy settings), our contribution here is a generalized approach that may be applied in any domain featuring social concerns and many characters. This method has also been used in the completed mixed-reality game *Bad News* (Ryan et al., 2016), as well as another game that we are developing called *Islanders* (Ryan, 2016).

The most related work that we are aware of is that of Méndez et al. (2014). Like ours, this

project targets a lightweight, generalized design and features a low-fidelity simulation procedure that evolves mutual character affinities over the course of many social interactions. In this system, character affinities vary according to a social exchange in which one character proposes an interaction to another character, who then chooses whether to accept or reject the proposal. If the recipient accepts, both of the characters' mutual affinities increase; upon rejection, the proffering character will lower her affinity toward the recipient. As such, we conceive of their system as a lower-fidelity version of *Comme il Faut* (CiF), the engine underpinning *Prom Week* (McCoy et al., 2014). While social exchanges in CiF have a rich array of preconditions and postconditions and are rendered in natural language, exchanges in Méndez et al. (2014) have simple conditions and are only represented symbolically. In sacrificing fidelity, their system gains in computational efficiency, which affords the generation of large social networks, a task for which CiF would likely prove too heavyweight.

Our project takes this lightweight approach further, with social exchanges that are even more abstract: interactions in our system are represented merely as functions that evolve affinities. While in Méndez et al. (2014), social networks of fifteen characters are generated, we have employed our approach to generate networks that span several hundred characters. Further, while their system includes only a single notion of affinity that subsumes both platonic and romantic feelings, we model these separately, as we discuss momentarily.

GENERATING CHARACTER SOCIAL NETWORKS

In this section, we outline our method for generating character social networks, which is a generalization of the approach that we employ in *Talk of the Town* (Ryan et al., 2015). Before outlining the simulation procedure that yields these networks, we define the core notions of our method and provide a listing of prerequisites that must be in place to implement this procedure in a game.

Core Notions

By our method, a social network is defined by the nonreciprocal feelings of friendship and romantic attraction held among a group of gameworld characters. Specifically, we rely on two core notions:

- **Charge.** A nonreciprocal *platonic affinity* held by one character toward another, represented by a scalar value. This value may be positive, signaling feelings of friendliness, or negative, indicating feelings of enmity. Broadly, charge increases as a character spends more time around someone who is compatible with her and decreases as she spends more time around a character who is incompatible with her.
- **Spark.** A nonreciprocal *romantic affinity* held by one character toward another, represented by a scalar value. This value may be positive, signaling romantic attraction, or negative, indicating romantic aversion; additionally, a default value (e.g., 0) may represent the absence of such feelings, in cases where romantic attraction is not applicable. As we explain below, spark evolves similarly to charge.

Prerequisites

The method that we present here requires several components to be in place for the simulation procedure to operate.

First, the method requires a handful of base components:

- **Basic modeling of time.** The method we present here operates by simulating the evolution of a social network over a series of timesteps. In *Talk of the Town*, we use day and night timesteps.
- **Basic modeling of space.** Each character must be located somewhere on a given timestep, which means that the space of the gameworld must be modeled in some way. We model space discretely in *Talk of the Town*, but space could just as well be modeled continuously.
- **Basic modeling of character personality.** Crucially, characters must be attributed personality models, and specifically a degree of *sociableness*. At points in the simulation procedure, a determination must be made as to the probability of a character engaging in a social interaction with a nearby character, and this requires knowing how sociable the characters are. In *Talk of the Town*, we use the famous *five-factor model* (Digman, 1990) (with each component represented as a floating-point value between -1.0 and 1.0), but any model may be used as long as a notion of character sociableness is included.

Additionally, the method requires a handful of higher-order notions that rely on the base components we have just listed. These are:

- **A notion of character proximity.** The simulation procedure must know which characters are nearby which other characters on a given timestep. In *Talk of the Town*, each character will be at a single home or business for the duration of a timestep, and so we say that characters are nearby one another on a timestep if they are located at the same home or business at that time. Using a continuous model of space, one might temporarily partition the space into discrete areas or else assert a threshold on proximity such that characters whose distance from one another does not exceed that threshold are considered nearby one another.
- **A notion of friendship compatibility.** As we discuss below, as characters spend more time together, their respective charge values for one another will increase or decrease according to a notion of their friendship compatibility, which is nonreciprocal. Specifically, this notion is used to determine a *charge increment*, which specifies how much a character's charge for another will increase (or decrease) in the case of a social interaction. As an example, in *Talk of the Town* we operationalize findings from the social sciences (Selfhout et al., 2010; Verbrugge, 1977) to produce a notion of friendship compatibility that captures the following (with regard to the five-factor personality model):
 1. *Characters with similar openness, extroversion, and agreeableness personality components are more likely to befriend one another.*
 2. *Characters with higher extroversion select more friends.*
 3. *Characters with higher agreeableness are selected as friends more frequently.*
 4. *Characters of the same gender are more likely to become friends.*

5. *Characters closer in age are more likely to become friends.*
6. *Characters of the same occupational status are more likely to become friends.*

Specifically, to calculate charge increments for two characters (relative to one another) in *Talk of the Town*, we do the following (numbers here indicate which finding from above each of these steps operationalize):

1. Calculate a base charge increment (for both characters) as the average difference between the characters' openness, extroversion, and agreeableness personality components, and then normalize this average to a scale between -1 (maximum disagreement) and 1 (maximum agreement).
2. Affect the charge increment held by each character according to the character's own extroversion personality component. If she is extroverted, this will boost the charge increment; if she is introverted, it will reduce the increment.
3. Affect the charge increment held by each character according to *the other character's* agreeableness personality component. If the other character is agreeable, this will boost the charge increment; if she is disagreeable, it will instead reduce the increment.
4. If the characters have different genders, reduce both of their charge increments.
5. Reduce both charge increments commensurately to the difference in age between the characters.
6. Reduce both charge increments commensurately to the difference in occupational status between the characters.

Again, we would like to emphasize that we only provide these details as an example, and that the method we present in this paper only requires that unidirectional charge increments be calculable for any pair of characters.

- **A notion of romantic attraction.** Like with charge, and as we discuss below, as characters spend more time together, their respective spark values for one another will increase or decrease according to a nonreciprocal notion of romantic attraction. Specifically, this notion is used to determine a *spark increment*, which specifies how much a character's spark for another will increase or decrease in the case of a social interaction. In developing such a notion for *Talk of the Town*, we again operationalize findings from the social sciences, in this case Luo and Zhang (2009). In the interest of saving space, we will refrain here from providing details of our implementation in *Talk of the Town*, but will emphasize that our method requires that unidirectional spark increments be calculable for any pair of characters (though in many cases these will be set to a value indicating non-applicability).

Finally, the method utilizes a crucial subroutine that must also be defined:

- **A procedure that places characters at locations on timesteps.** Rather than placing characters randomly each timestep, characters should decide where to go according to concerns that are at play in the gameworld. In *Talk of the Town*, characters carry out daily routines by which they may go to work, go on errands, visit neighbors or friends or family members, visit places of leisure (*e.g.*, a restaurant), or stay at home.

This works well with the method we describe in this paper, because the routines ensure that characters develop relationships with characters that they would believably spend time around (given the game's diegesis), *e.g.*, co-workers or neighbors in the case of *Talk of the Town*.

Simulation Procedure

Having these required components in place, the simulation procedure itself is quite simple. In this section, we will explain this procedure by providing pseudocode, with commentary appearing as necessary. In the next section, we outline optional extensions to the procedure.

while world generation is not over

In *Talk of the Town*, our procedure terminates when a specified timestep is reached, which represents the time at which gameplay takes place.

advance one timestep

for each character in the gameworld

decay charge and spark toward all other characters

Specifically, decay the charge and spark values that this character holds for all other characters in the gameworld by regressing them toward 0. (The degree of decay is a configurable parameter.) This makes strong feelings of friendship or enmity (or romantic attraction or aversion) weaken as time passes without the character interacting with the subject of such feelings. (Below we note an extension to this procedure that decays charge and spark differently.)

place character at a location

As noted above, in *Talk of the Town* this step relies on a notion of character daily routines.

for each location in the gameworld

determine which characters at location will interact

For believability, this should be done using the components of the characters' personality models that specify how sociable they are. In *Talk of the Town*, this step also considers the charge and spark values that characters hold for one another, which triggers rich feedback loops that may reinforce emerging friendships. That is: characters are more likely to interact with better friends, and interactions between friends will cause friendships to be strengthened; in turn, such strengthening will increase the likelihood of future interactions between the friends, which will work to further strengthen the friendships, and so forth. This is how *best friend* relationships emerge in *Talk of the Town*.

for each dyad of characters that will interact

for each character in the dyad

update charge and spark toward other character

Specifically, update these values by the appropriate charge and spark increments.

Optional Extensions

In this section, we outline a number of optional extensions to our method, some of which we have ourselves utilized in *Talk of the Town*.

- **Evolving different kinds of affinity.** As a fundamental type of extension, applications of this method could evolve types of affinity beyond platonic and romantic feelings, such as *reputation*, for instance. To do this, one simply needs to define a way to determine the increment for such affinity, given any two characters. Then, affinities of this type evolve just as charge and spark do in the way we have explained.
- **Instantiating discrete relationships.** A basic extension to our method involves instantiating discrete relationships as affinity thresholds are reached. In *Talk of the Town*, we instantiate an *acquaintance* relationship whenever the absolute value of a character's charge toward another character is less than a specified threshold, and *friendship* and *enmity* relationships when such values eclipse specified positive and negative thresholds, respectively. For romantic relationships, threshold checking can be used as an abstract simulation scheme, or else thresholds can be used to trigger romantic gestures between characters, such as an invitation to begin dating. We also instantiate *best friend* and *love interest* relationships by querying for characters' maximum charge and spark values, respectively.
- **The interpersonal circumplex.** Once affinities have been evolved, they can easily be combined with character *status* relations (Johnstone, 2012) to produce the *interpersonal circumplex* (Wiggins, 1996), a powerful model of interpersonal relations that maps affinity against patterns of dominance and submission. In *Talk of the Town*, we have a number of reified status relations: boss–employee, senior employee–junior employee, elder family member–younger family member.
- **Recording data about social interactions.** Because our method crucially relies on a procedure that believably simulates when and where character social interactions occur, a simple extension with potentially large payoff involves the tracking of data about the time and location of social interactions. In *Talk of the Town*, we use such data to instantiate interesting character beliefs that support intrigue surrounding the larger mystery that underpins gameplay (Ryan et al., 2015). While our approach abstracts away the *reasons* for relationships being in certain states, this data could be used along with circumplex models to *retrofit* histories of discrete social exchanges.
- **More nuanced affinity decay.** While we specified in the above description of our central simulation procedure a basic approach to affinity decay, in *Talk of the Town* we employ a more nuanced approach that decays platonic and romantic affinities in different ways: we decay platonic affinities by reducing the charge characters have for one another as time passes without them interacting (as in the description above), but we decay romantic affinities by reducing the spark *increment* a character has for another as the two spend more time together. The latter type of decay supports romantic

attraction that evolves quickly but eventually converges, as well as romantic obsessions that persist even without extended interaction with the target of such affection (both of which are media tropes we consciously operationalize).

- **Extrapolation.** Rather than simulate every timestep during world generation, in *Talk of the Town* we simulate a handful of timesteps a game year and extrapolate charge and spark increases according to the amount of time since the last simulated timestep.

DISCUSSION

Through extensive testing of this method (and its use in *Bad News*, as mentioned above), we have found that the social networks that it generates exhibit interesting phenomena—love triangles, unrequited love, asymmetric friendships, family feuds, industry rivals—without requiring retroactive modification or inspection for failure cases. Our approach may be used in any game that features many characters, but it is only intended as a lightweight method for producing a *base* social world. To make full use of the method, authors should consider the optional extensions to it that we enumerated above. This way, any particular character social concerns that are important to the project may be amplified accordingly.

BIBLIOGRAPHY

- Adams, T. (2015). Simulation principles from Dwarf Fortress. In S. Rabin (Ed.), *Game AI Pro 2: Collected Wisdom of Game AI Professionals*.
- Afonso, N. and R. Prada (2008). Agents that relate: Improving the social believability of non-player characters in role-playing games. In *Entertainment Computing*.
- Digman, J. M. (1990). Personality structure: Emergence of the five-factor model. *Annual Review of Psychology*.
- Evans, R. and E. Short (2014). Versu—a simulationist storytelling system. *Computational Intelligence and AI in Games*.
- Garda, M. B. (2013). Neo-rogue and the essence of roguelikeness. *Homo Ludens*.
- Johnstone, K. (2012). *Impro: Improvisation and the theatre*. Routledge.
- Luo, S. and G. Zhang (2009). What leads to romantic attraction: Similarity, reciprocity, security, or beauty? evidence from a speed-dating study. *Journal of Personality*.
- McCoy, J., M. Treanor, B. Samuel, A. A. Reed, M. Mateas, and N. Wardrip-Fruin (2014). Social story worlds with *comme il faut*. *Computational Intelligence and AI in Games*.
- Méndez, G., P. Gervás, and C. León (2014). A model of character affinity for agent-based story generation. In *Knowledge, Information and Creativity Support Systems*.
- Ryan, J. (2016). Diegetically grounded evolution of gameworld languages. In *Proc. Procedural Content Generation*.
- Ryan, J. O., A. Summerville, M. Mateas, and N. Wardrip-Fruin (2015). Toward characters who observe, tell, misremember, and lie. In *Proc. Experimental AI in Games*.
- Ryan, J. O., A. J. Summerville, and B. Samuel (2016). *Bad News: A game of death and communication*. In *Proc. CHI*.
- Selfhout, M. et al. (2010). Emerging late adolescent friendship networks and big five personality traits: A social network approach. *Journal of Personality*.
- Verbrugge, L. M. (1977). The structure of adult friendship choices. *Social Forces*.
- Wiggins, J. S. (1996). An informal history of the interpersonal circumplex tradition. *Personality Assessment* 66(2).