

The Affinity Project

WFC


May 2022

Affinity

- A program written in Java, using JavaFX for graphical animation.
- Simulates a community of interacting social entities.
- The simulation is intended to model how social relationships develop through interaction with others.
- The simulation is intended to model the use and development of *personal concerns*: Affinities, attitudes, beliefs, values, opinions
- An entity may use personal states about itself, and about the others with which it interacts, including imagined others.
- A computational model of how entities reason about social concerns.
- Intended to model aspects of personhood and to explore ideas of spiritual intelligence involving social relationships.
- The implementation is incomplete: It may not model your favourite thing.

The World

- The world is the interior of a box, in which physical laws of motion are simulated.
- Entities are represented as coloured spheres.
- Entities may raise and lower one of several different visible signs.
- Entities can develop friendships, and friendships are shown by a visible link.
- Entities can seek care and offer care to other entities.



In this screenshot, a few friendships have been made. Some entities have raised their 'help needed' sign or 'giving care' sign.


User Interface

- The Graphical User Interface (GUI) controls the simulation.
- Information about the population as a whole is presented by the GUI.



Mimic Diagram

- An entity can be selected, and its status displayed on a mimic diagram.
- The behaviour of entities is governed by a set of its own *narratives* that are processed concurrently by each entity.




Mimic Diagram

Tick, friends, energy

Concerns

Mission Status

Narratives



Concerns

- **Basic Values** (Schwartz): SELFDIRECTION, STIMULATION, HEDONISM, ACHIEVEMENT, POWER, SECURITY, CONFORMITY, TRADITION, BENEVOLENCE, UNIVERSALISM.
- **Attitudes**: FORGIVING, CREDIBLE, SCEPTICAL, ...
- **Social Objects**: OPINION, COMPAT, ...

Basic Values

- SELFDIRECTION, STIMULATION, HEDONISM, ACHIEVEMENT, POWER, SECURITY, CONFORMITY, TRADITION, BENEVOLENCE, UNIVERSALISM
- Of these, only BENEVOLENCE and UNIVERSALISM are currently implemented.
- BENEVOLENCE governs an entity's care for its in-group.
- UNIVERSALISM governs an entity's care for its out-group.
- Most of the others will be implemented in due course.

Attitudes

- Examples: FORGIVING, CREDIBLE, SCEPTICAL
- Currently only FORGIVING is implemented, and governs the reconciliation process, when entities meet ex-friends and decide whether to re-friend.
- CREDIBLE, SCEPTICAL will be implemented next, to do with how an entity processes the opinions of others.

(Note the official statement of St Albans Diocese is phrased as attitudes: Generosity, Joy, Imagination, Courage)

Social Objects

- Social objects are the *topics* about which entities may have opinions, affiliations, and preferences.
- Entities may also have opinions, affiliations, and preferences about what social object state they are compatible with in others.
- Currently the only topic implemented is FAVOURITE COLOUR.
- Entities can have an OPINION about a FAVOURITE COLOUR, and they can have a COMPAT, which is an opinion about the possible FAVOURITE COLOUR of others.
- Compatibility between social object states is one factor in forming relationships between entities.
- In the simulation, an entity's body is displayed in its favourite colour.

States

- An instance of a particular concern is called a **state**.
- A state consists of a triple: concern, degree, affinity.

$\langle \text{BENEVOLENCE}, N(0.6, 0.2), N(0.8, 0.01) \rangle$

Concern


Degree 



Affinity 


In this model, all quantities are numbers between 0 and 1. This comes from a simple conjecture that brains process non-negative values that can be modelled as probabilities. When a quantity is needed, it is randomly sampled from a Normal distribution. This introduces some stochasticity.

- The degree is the 'amount' of the concern held by the entity.
- The affinity is how much it is held important for action (motive force).
- The 'brain' of an entity consists of a collection of states and narratives.

Compatibility

- The compatibility between two entities is defined in terms of social object states. Because state quantities are Normal distributions, the z-score can be used to judge compatibility between states.
- Compatibility is defined on a 4-point qualitative scale: Similar, Marginal, Different, Separated.



Issues with a compatibility model

- Compatibility can be multi-dimensional: Several different social object states can be involved in compatibility judgements, and states can have multiple components (e.g. the 3-component RGB colour model is used).
- A simple model is used, in which distributions are averaged, and the z-score is calculated between average distributions.
- The simple model can be improved. In everyday life, so-called 'red flags' are when an incompatibility is recognised in one dimension that may overrule an otherwise compatible comparison.
- Future work should look into a better compatibility model.

Relationships

- This work starts from the observation of MacDorman and Cowley (2016) that a benchmark of personhood is engagement in long-term relationships.
- Currently, the model implements 'friendship' according to simple criteria. Friendships can be made and broken according to proximity, compatibility with social objects, and certain other concerns. Attentive to Dunbar's work.
- Currently there is only one level of friendship.
- A richer range of relationships will be implemented, such as influencer-follower, membership groups, relationships based on complex concerns (e.g. 'blood is thicker than water').

Initial conditions

At the start of the simulation, each entity in the population is initialised with a random position, velocity, and energy; a set of concern states with randomly set degrees and affinities, and a set of six narratives.

The concern states are

- SocialObject.Type.**OPINION** *(The only topic so far is favourite colours)*
- SocialObject.Type.**COMPAT** *(The only topic so far is favourite colours)*
- BasicValue.Type.**BENEVOLENCE**
- BasicValue.Type.**UNIVERSALISM**
- Attitude.Type.**FORGIVING**)

The Energy Economy

- The 'health' of an entity depends on how much energy it has.
- If the energy of an entity drops too low for too long, it dies.
- The life span of an entity therefore depends on the rules set by the energy economy.

EFFICIENCY = 0.005;	<i>// how much energy is required by entity per unit of velocity</i>
SOLAR_ENERGY = 0.001;	<i>// given if Sun out, consumed if Sun in</i>
HAVINGFRIENDS = 0.01;	<i>// times maximum affection over all friends</i>
NOFRIENDS = 0.01;	<i>// loses this much per tick if no friends</i>
LIMIT_ENERGY = 0.5;	<i>// care is sought if energy level < 0.5</i>
CARE_ENERGY = 0.05;	<i>// Care causes increase of energy in units per tick</i>
	<i>// This is > CARELOSS because the environment adds energy exogenously</i>
HEALTHY = 1.5;	<i>// Considered healthy (the gap with NEEDCARE is for hysteresis)</i>
CAREGIVER = 2.0;	<i>// Caregivers need to have at least this before responding</i>
CARESTOPS = 1.0;	<i>// Caregivers stop giving care if their own energy < this</i>
CARELOSS = 0.01;	<i>// Caregivers lose this per tick</i>

Other Concerns -- Beliefs, Norms, Traits

- **Beliefs** are ideas about how true it is that things are related in particular ways.
- Examples of beliefs are: War never solves problems, Malta is larger than France, the Earth is a flat disc, and theologians are wise.
- Beliefs vary in how certain *we* are that they are true.
- General beliefs that people hold about how the world functions are called *social axioms* (Leung & Bond, 2004).
- Unlike values, beliefs refer to the *subjective* probability that a relationship is true, not to the importance of goals as guiding principles in life.

Adapted from Schwartz (2011)

Other Concerns -- Beliefs, Norms, Traits

- **Norms** are standards or rules that tell members of a group or society how they should behave.
- Examples of norms are “children should be seen and not heard,” “we should stand up when the national anthem is played.”
- Norms vary on a scale of how much we agree or disagree that people should act in a specific way. Our values affect whether we accept or reject particular norms.
- Because norms prescribe behaviours with specific *consequences*, we are more or less inclined to accept them depending on whether these consequences are compatible or in conflict with our valued goals.
- More generally, because norms are *social expectations*, we are more or less inclined to accept them depending on how important conformity vs. self-direction values are to us.

Adapted from Schwartz (2011)

Other Concerns -- Beliefs, Norms, Traits

- **Traits** are tendencies to show consistent patterns of thought, feelings, and actions across time and situations.
- The same term can refer to both a value and a trait (e.g., wisdom, obedience). However, people who exhibit a trait may not value the corresponding goal and those lacking a trait may value the corresponding goal highly.
- For example, a person may behave obediently yet not value obedience and behave foolishly yet value wisdom highly.
- Traits vary in the frequency and intensity with which people exhibit them.
- They describe what people are like rather than what people consider important. People believe their values are desirable but may consider their traits positive or negative.

Adapted from Schwartz (2011)

Narratives

- The behaviour processes of an entity are represented by *narratives*.
- A narrative is an ordered sequence of *steps* that defines an information processing behaviour extended in time.
- A narrative has a first step and a last step. Performance of a narrative begins at the first step. Each step is a condition-action rule. If the condition is satisfied, the action is executed. If the condition is not satisfied, then nothing happens: the condition is tested again at the next opportunity. A narrative is discarded after the last step.
- Narratives are performed concurrently.
- A narrative also has a small amount of local state, and some of the steps of a narrative can be performed in a particular mode called a *mission* (defined later).
- Currently each entity is equipped with six different narratives.

Step

- Each step of a narrative is a condition-action rule represented as a triple:

⟨ OPERATION, CONDITION, ACTION ⟩

There are four operations with the following semantics:

- IF *condition action*. If the *condition* tests true, then perform the *action*.
- WAIT condition action. Wait until the *condition* tests true, then perform the *action*.
- WHILE condition action. Repeat, while the action tests true, perform the action.
- LOOP. Go back to the beginning of the narrative.

Six Basic Narratives

1. Motility

IF FINDCENTRECOMPATFRIENDS MOVECLOSERPOINT

IF FINDCENTREFRIENDS MOVECLOSERPOINT

IF FINDNEARCOMPAT MOVECLOSERENTITY

LOOP NOCONDITION NOACTION

2. Maintain Friendships

IF NOCONDITION MAINTAINFRIENDS

LOOP NOCONDITION NOACTION

Six Basic Narratives

3. Make Friends

WAIT MEETPROXIMALCOMPAT GAINFRIEND
LOOP NOCONDITION NOACTION

4. Ask Help

WAIT NEEDCARE RAISENEEDSIGN
WAIT AMHEALTHY LOWERNEEDSIGN
LOOP NOCONDITION NOACTION

Six Basic Narratives

5. Give Friends Care

```
WAIT FINDNEEDYFRIEND STARTCAREMISSION  
IF NOCONDITION RAISECARESIGN  
WHILE CAREMISSIONFAR MOVETOMISSION  
WHILE CAREMISSIONNEAR GIVECARE  
IF NOCONDITION ENDMISSION  
LOOP NOCONDITION NOACTION
```

Six Basic Narratives

6. Give Others Care

WAIT FINDNEEDYOTHER STARTCAREMISSION
IF NOCONDITION RAISECARESIGN
WHILE CAREMISSIONFAR MOVETOMISSION
WHILE CAREMISSIONNEAR GIVECARE
IF NOCONDITION ENDMISSION
LOOP NOCONDITION NOACTION

Governors

- A Step of a narrative is processed depending upon the current state of the entity and a set of (learnable) weights.
- For example, the FINDNEEDYFRIEND action is executed only if the entity has sufficient energy to be a care giver and has a certain amount of benevolence. This is implemented in the program as the Java condition

`(!e.onMission(Mission.Objective.GIVECARE)`

`&& e.body.energy > Energy.CAREGIVER`

`&& e.sampleValue(BasicValue.Type.BENEVOLENCE) > 0.6)`

Note added in proof: In principle, governors could be implemented using a Step condition

Missions

- A narrative may start and end a *mission*. A mission is defined by an *objective*, a *target*, and a *status*.
- When a mission is started, each of the steps following in the narrative is tested with an *objective* (a type of conditional) before it is executed. If the objective tests true, then the step is taken. If the objective tests false, the mission is cancelled: All the steps before ENDMISSION are skipped.
- A mission has the status {LIVE, CANCELLED, COMPLETED, CLOSED} .
- Missions are useful because they maintain a set of conditions over a sequence of Steps. This is more powerful than a traditional **do..while** because the condition is tested at every step.
- Conjecture: Self-awareness is about missions.

Example mission

Give Friends Care

WAIT FINDNEEDYFRIEND STARTCAREMISSION

IF NOCONDITION RAISECARESIGN

WHILE CAREMISSIONFAR MOVETOMISSION

WHILE CAREMISSIONNEAR GIVECARE

IF NOCONDITION ENDMISSION

LOOP NOCONDITION NOACTION


During these steps, the following mission state is held by the narrative:

Target: the entity found

Objective: Does target need care

Status: LIVE

A mimic diagram displays the current state of the narratives, and which step is currently being processed.



Every narrative of every entity is progressed concurrently on each clock tick (about 100ms). This is for user convenience only: No commitment to synchronous operation is implied.

Learning

- Not yet implemented.
- The plan is to implement adjustment of governing weights so the community can develop over time.
- One easy way is to assay the concern states of the entities that die and the entities that flourish, and use this as an *online auto-supervised* training set to adjust the concern states of future populations.
- Another way is to use matrix factorisation to determine the latent states of successful and unsuccessful groups.
- Conjecture: Learning results in entities becoming more benevolent and caregiving, because the energy economy rewards long term friendships and receiving care.

Further work

- Implement more concerns: Social objects, Basic values, Attitudes, Beliefs, Norms, Traits.
- Implement more topics for social objects plus a better model for compatibility judgements.
- (Re)implement death and loss.
- Implement learning, through online auto-supervised learning, and through factorisation into latent features.
- Implement a greater range of relationships.
- Implement the narratives around opinion forming, belief, and influencing. Previous work on modelling human opinion forming on social media can inform this.