neighbourhoods

Distributed Social Sensemaking

Primitives for Groupware Design

Pablo Somonte Ruano, Emaline Friedman & Jean Russell
with contributions from the Neighborhoods team

July 2022

Table of Contents

Introduction

Social sensemaking is *about* the social world. It is a process of orienting to people, spaces, scenarios, and representations. For this orientation to occur, the latter need to be defined in some way. Social sensemaking is also a social process in itself. It's articulable *between* members of a social group. Individuals use their goals, aims, and desires, conscious or not, to make sense of the social world on their own. For social sensemaking to be a group rather than an individual activity, the group must have an apparent, commonly defined set of interests, goals, and/or values.

To support the proliferation of digital groups with specific cultures in their communities, in Neighbourhoods, *community activators* decide on the collective sense-making mechanics that formalize their group's norms, rules, and culture. We think of community activators as the starters and/or stewards of a community, who can gradually distribute their initiative through the group as part of collective processes of sensemaking, and eventually governance.

Probably our biggest gripe about platforms is that, in trying to "do it all", they end up imposing the cultural norms of developers. Propped up by sleek, clean interfaces, platforms project a false universality that leaves almost no room for communities to practice their own cultures. To the degree this is possible (always a lively debate!) we believe in a separation between generic organizing tools and the cultural components that make them come alive for people.

Neighbourhoods tools aim to support social sensemaking by *augmenting* the unique ways people come together into social bodies, rather than diminishing these. The term "sense-making" reflects our intention to amplify the ethnographic rather than the mechanistic.

We understand sensemaking as any individual or shared understanding of resources and agent's capabilities within groups derived from the aggregation of social data and the subsequent computations performed on that data.

These processes specifically impact:

- an agent's access to groups
- their rights inside those groups
- the visibility of resources in specific contexts

Because this is all happening via distributed, p2p software, social sense-making as we understand it implicitly requires modularity, transparency of computations that are normally hidden in a "black box", cultural autonomy, and community self-governance.

Thus, we understand group processes such as governance, reputation, coordination, feed-ordering, reactions, ratings as *manifestations of sensemaking.* Rather than rushing to create processes for these higher level operations, Neighbourhoods is focused on what is required to have these things be useful in the first place.

Distributed social spaces

Every neighbourhoods is comprised of *groupware:* Any software that groups use to connect with each other, share information, ideas, memories, and goods, and coordinate their actions. In any given neighbourhood, groupware "modules" are lightweight Holochain hApps.

Social sense-making is social in two ways. First, it is *about* the social world. It is a process of orienting to people, spaces, scenarios, and representations. For this orientation to occur, the latter need to be defined in some way.

Second, it is a social process, which is to say, articulable and/or visible *between* members of a social group. Individuals use their goals, aims, and desires, conscious or not, to make sense of the social world on their own. For social sense-making to be a group rather than

an individual activity, the group must have an apparent, articulable set of common interests, goals, or values.

Therefore, when groups pursue sense-making as a common activity, these interests, goals, and values are cultivated, becoming more refined and more actionable in the process. Social Data helps make sense of social interactions and shared experiences.

> It's not necessarily the case that neighbourhoods are "social media" — even though our MVP v.1, built for demonstrative purposes, was a meme network.
>
> However, they are spaces for the social coordination of agents and resources, irrespective of whether these are physical, digital, or media artefacts. For neighbourhoods to effectively coordinate agents and resources, then, social sense-making is critical!

How it's done on the centralized web

The prize of Facebook, that which allows it to sell micro-targeted ads, to enable government backdoors, and to develop new products on the basis of user data is called the social graph.

The social graph is what enables you to see your "friends in common", and building it out right before its IPO is why so many new input options (more moods, more genders, more tagging options, more 'places') were introduced. The social graph is the money maker for monolithic platforms, which means we're accustomed to a glut of social data for this purpose rather than for sense-making aimed at enabling community self-governance.

In distributed environments that do away with profiting off the sale of user data , data can become *socialized* — a set of information used for and by the community that produces it. This requires a rethink, and the concepts and categories below should help you do just that.

Resources, Agents and Methods

Sensemaking is set up by specifying the **resources** the neighbourhood forms around. Are these images? Movies? Articles? Crowdfunds? Another way of asking this is: “what is the *content* in the neighbourhood”? **Agents** are the “who” that perform actions, make contributions, curate, and seek information. **Methods** involve the way math is applied to social data to derive insights and new information. Using Neighbourhoods, these computations are serverless and also redundant.

Contexts, Access and Visibility

Once resources, agents, and methods are specified, output computations impact the **access** of agents (permissions) and the **visibility** of content (such as a feed order or content listing). These two processes hinge on when and how resources, agents, and sensemaking procedures become legible to participants.

Neighbouroods contain many different spatial visualizations that operate through unique rights, permissions, and access points. We call these **contexts.** Access to resources depends on social linking between agents. Access to contexts depends on role credentials, and access to neighbourhoods themselves depend on membership credentials. These are important levers in creating and maintaining neighbourhoods — rules for how **crossing these membranes** works is a large part of initial configuration.

In the following sections, we’ll elucidate factors involved in social sensemaking with Neighbourhoods. Coincidentally, we’ve been thinking about what a neighbourhood for soliciting, reviewing, and accepting research for publication would look like.

So, everything you want to know about sensemaking, and our standards for social data, will come to life by looking at this use case!

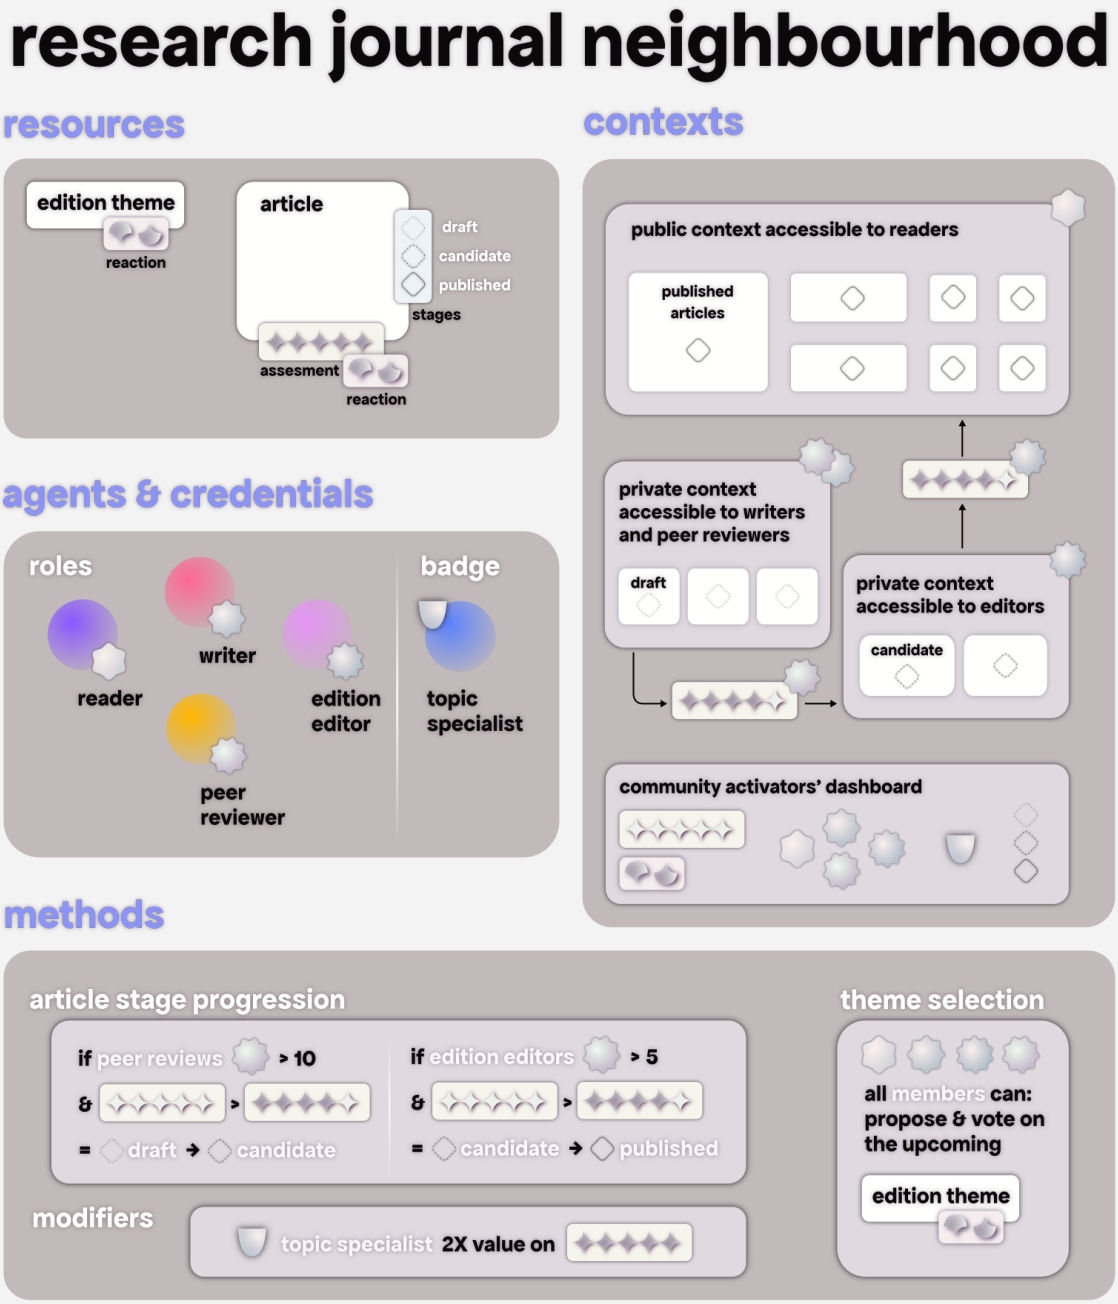

We hope our ontology for culture design is thought-provoking, yet grounding. As always, we want to support community activators in developing a sophisticated approach to creating neighbourhoods that embody their specific culture. We also hope to enable developers in taking advantage of the affordances of Neighbourhoods' Social Sensemaker tooling.

Social Network Primitives

The “social network primitives” and their interconnections represent a key, preliminary step toward our design goals for neighbourhoods. These are:

- modularity
- composability
- Interoperability

Resources, contexts, agents and “object relations” (including credentials, social linking and thresholds) — are at the root of our evolving understanding of social sensemaking. To make it all more tangible, we’ll explain these concepts through a prototypical **research journal neighbourhood** — an example neighbourhood meant to convey user-facing features.

Resources

Resources are the elements within a neighbourhood that are subject to sensemaking operations. Tweets on twitter, posts on Instagram, movies in Rotten Tomatoes or publications in Medium are all examples of resources. A resource can be composed of multiple elements.

For example, a ‘post’ on Instagram is a resource composed of:

- Between 1 and 5 images, or videos that are shorter than a minute long
- A text caption
- One or more agents credited as authors
- None to multiple agents tagged inside each picture or video
- A comment section (which in itself contains a distinct sub-resource: the comments composed of texts)

What is considered to be a resource largely determines the granularity of sensemaking that can be performed upon it. In the case of the Instagram post, one could imagine more granular sensemaking if one could react to each image or video *within* the post rather than all of these elements taken, together, to be one single resource.

In our research journal example, the main resources are the ‘articles’ which are subject to sensemaking through ‘assessments’ by members of the neighbourhood who hold specific roles. Additionally, articles go through three distinct stages ‘drafts’, ‘candidates’ and ‘published’. (More information on assessments, reactions and roles coming later).

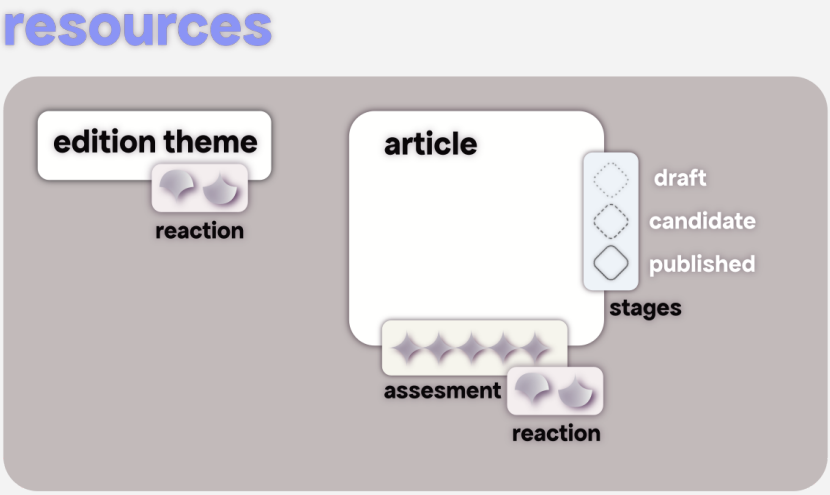

An example of a resource that is not a ‘post’ or what we’d normally consider ‘content’ is the edition theme in our research journal example. In the journal NH, a theme can be proposed by any member, and other members can react to the proposal with a reaction in the form of upvotes and downvotes to state their preference for the theme of the upcoming edition of the journal. Upvoting and downvoting proposed edition themes exemplifies how reactions operate in a “proto-governance” capacity. They are useful in carrying out simple collective choice operations, and one could easily do the same with more nuance, like time-limited or numerically-scarce reactions per member. In this way, making a decision in a neighbourhood becomes part of a broader spectrum of sense-making.

Contexts

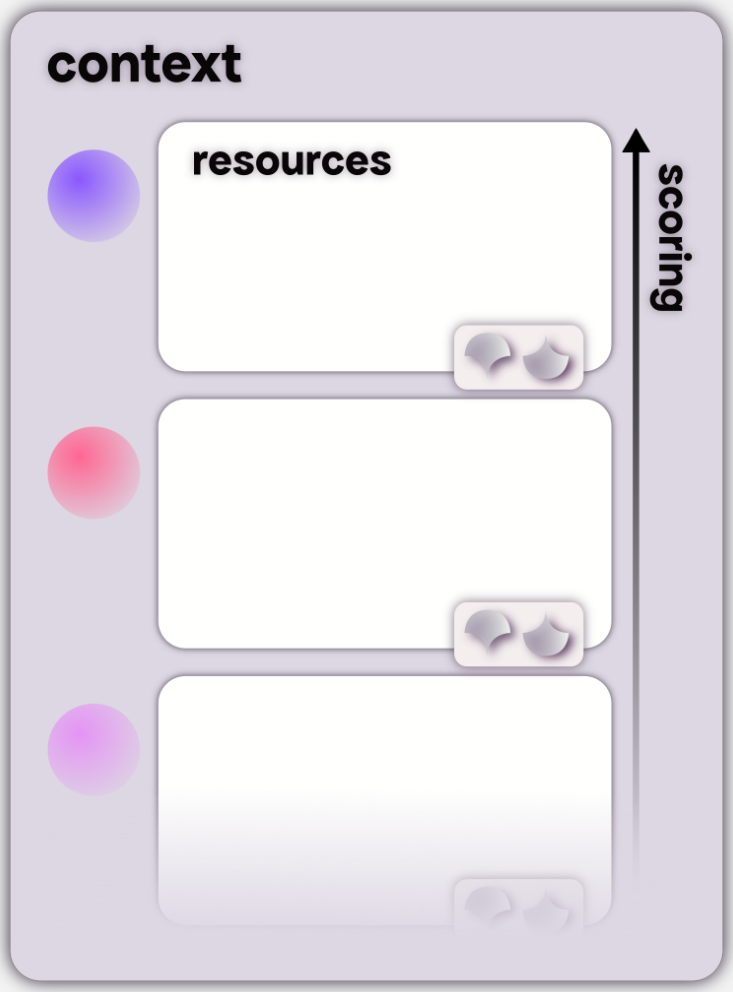

Contexts are the spaces *populated* by resources. Twitter's feed, Instagram's profile grid, and a list of houses on Airbnb are all examples of contexts in current social networks. Contexts, as the name implies, contextualize resources creating meaning by ordering them in specific ways. Twitter's feed can be ordered according to its black-boxed algorithm, or chronologically, while Instagram's profile grid is strictly chronological. On Airbnb, context emerges as the result of a search and houses are ordered according to the reviews of renters combined with paid sponsorships by sellers.

How resources are ordered in a context is fundamental for processes of collective sense-making, where values and culture are expressed in specific ways.

Social applications today have a variety of contexts shown in a UI at the same time. In Instagram's home page there's a feed context that shows the posts of people you follow, and there's also a horizontal context with a list of profiles that takes people directly to a 'stories' context. There are also sub-contexts that depend on other components — each Instagram post has the sub-context of a 'comment-section' with its own sub-resource the 'comments' which is subject to its own algorithmic ordering. In other cases, multiple contexts show the same resource, at the same time, but with different ordering. For example, in a live-stream context one could find a context showing comments in a

chronological order, in addition to a context showing the comments of those who have donated money or who are special members of the community.

In our ontology, **neighbourhoods are collections of contexts**, with their own particular algorithms determining the ordering and visibility of resources. Contexts are configurable to be accessible to only certain groups or to all members, which creates rich, “multi-membraned” social spaces.

The ‘edition theme’ resource in our example is propelled by upvotes and downvotes that move themes to the top at the end of a specified period of time. The theme with the highest ratio of upvotes to downvotes becomes the resource at the top of the context, making it the theme for the upcoming edition and setting in motion the next round of submissions and reviews.

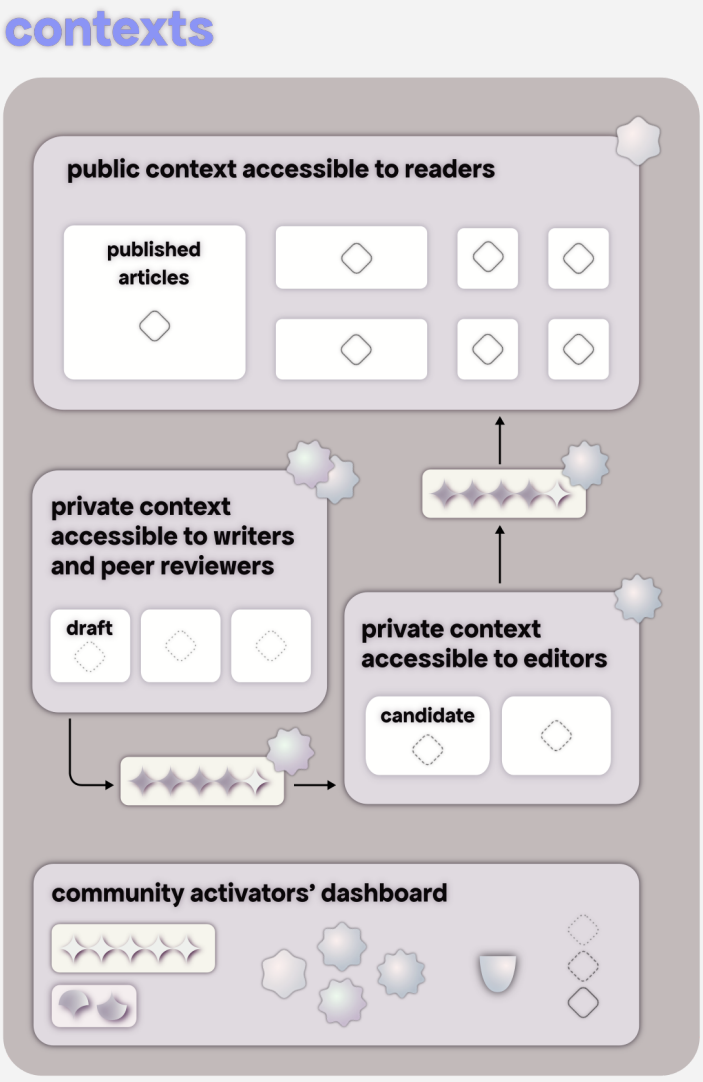

Each contexts in the research journal neighbourhood benefits from a variety of scoring algorithms, depending on who accesses them. Peer reviewers might want to see the articles with fewer reviews first, or those that coincide with their topic of expertise. Contributors, on the other hand, may want to see the best reviewed articles first in order to get a sense of the quality standard of the community.

A context for published articles might resemble the front page of a newspaper, where the best rated article occupies more visual space in the UI, and could include some of the written comments that accompany the reviews made by the editors of the journal. Subsequent articles would then occupy progressively less space according to their review. But other, qualitative scoring algorithms could also come in handy here! For example, an algorithm showing resources grouped by topics that your friends, or those with whom your share interests, have engaged with that you have not yet encountered.

Agents

We adopt the language of “agents”, rather than “users”, in keeping with Holochain terminology. It’s defined here in the glossary as anyone or anything acting with agency, such as a human or a bot, where agency is simply the ability to act in a specified environment. For the purposes of distributed computing, agents can also have ID(s) utilizing public key cryptography. From a notion of agents, it is possible to create different forms of membership, rights, access points, and neighbourhood- or context-specific roles. The next section covers how our research journal examples create complex permissions and member types through ‘credentials’.

Object Relations

Object relations are about the interconnections between different components of the neighbourhoods ecosystem, and how those relationships are mediated. The main impact of object relations is determining access to, and rights over, other components.

We identify three main object relations: **credentials**, **social linking** and **thresholds**.

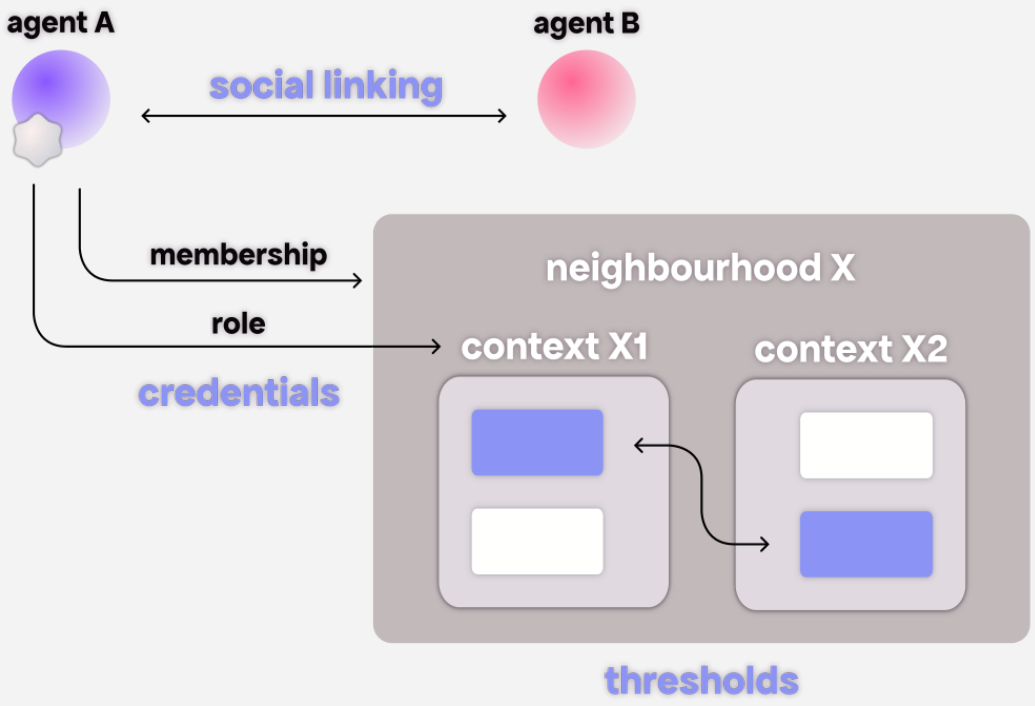

Credentials

Credentials refer to the relations between an agent and a collective space, informing access to those spaces and the rights within them. In our ontology, we refer to the relationship between an agent and a neighbourhood as a ‘membership’ credential (answering, e.g. “Can I enter this neighbourhood?”). The relationship between an agent and a context *within* that neighbourhood we call a ‘role’ credential (determining, e.g. “What can I do within this neighbourhood?”).

In the research journal neighbourhood, becoming a member grants you a reader role, enabling access to the ‘published’ articles. To become a contributor or a peer reviewer, one needs an invitation from the ‘edition editors’. While both the ‘contributor’ role and the ‘peer reviewer’ role give access to the same context where ‘drafts’ are visible, the rights of each role are different. Only contributors can submit articles as drafts, and only peer reviewers can assess them.

Generally, credentials also impact how a member's contributions are valued. Through 'awards', 'achievements' and 'badges', members receive credentials (automatically via computations or manually via direct grant by other members) that modify their rights. In this example, a badge signifying 'expertise in a topic' is awarded by editors to other members. This badge applies a 2x multiplier to the value of reactions and assessments of the holder when they interact with an article that has a tag coinciding with their area of expertise. In this way, the neighbourhood signals acknowledgment of the holder's opinion and formalizes this acknowledgment through a credential and the impact the credential has on the group's sensemaking capacity.

Credentials manifest one's reputation (and the visibility and authority it entails) within a certain group, purely on the basis of the group's formal rules about valuing activity and contribution.

Social Linking

Social linking references agents' connections with one another. On social media platforms, Facebook "friends" and Twitter "followers" impacts the information we see, but also how much personal information is revealed or how easy it is to contact someone directly. By the same token, social linking enables limitations on these connections, as when one member mutes or blocks another. In a neighbourhood, social linking might be used to change the order of resources to reflect one's social-interactional graph, like showing the posts of friends first or creating posts that are restricted only to close friends.

Interestingly, in the p2p world of Holochain social linking plays a critical role in maintaining the technical integrity of the network; where density of social links is tied to the monetization of surveillance on today's social media platforms, and to whole-network verification and gas fees on emerging blockchain-based socials,

> social links on Holochain pertain to whose information one replicates, in order to produce sufficient redundancy for network resilience, validation for network consistency, and security.

In neighbourhoods, as with Holochain more generally, social linking is genuinely techno-social — a microcosm of the technical trust being a "full node" of a blockchain once indicated, plus the social consequences of interdependence, implicit in the collective construction of social networks.

Thresholds

Finally, thresholds express the relationship between resources and contexts. They determine whether a specific resource has access to a specific context, and what members can do with it. In our example, there are 'stages' for the 'articles' resource that act as distinct thresholds that determine if an article is at a 'draft', 'candidate' or 'finished'

state. What state the article is in determines what roles have access to them and what interactions are possible, so a member with a ‘readers’ role will only have access to the context where ‘finished’ articles are shown, while members with the ‘contributors’ role can submit new articles as ‘drafts’ and see that context but only those with the ‘peer reviewer’ role can provide assessments to those drafts. Finally, only those with the ‘edition editor’ role can assess the articles in the ‘candidate’ state to bring them towards a ‘finished’ state.

It’s noteworthy here that the dimension of time plays a considerable role in determining various thresholds. Time can be considered in terms of periodicity (how often over what period), event time (events might trigger certain actions, regardless of clock time), and clock time (objective amounts of time elapsed before some event occurs or some threshold may be crossed). In this example, drafts need to have at least 10 assessments by peer reviewers and an average of 4 or more stars for them to automatically transition into the candidate state, there they need at least 5 assessments by edition editors and an average of 4 or more stars to automatically transition into ‘finished’ state where they will be visible to all readers.

Sense-making Data Types

Sense-making data types are the primitive components of the Neighbourhoods ecosystem; the common language that makes it possible for radically different communities to arrive at shared understanding when necessary. At the level of individual neighbourhoods, these data types allow agents to interact with resources through small gestures that attach subjective meaning to them.

We encounter the web2 equivalent of sense-making data types in various forms today such as ‘likes’ on Twitter, emoji reactions on Facebook posts and reviews and star ratings on Amazon. Other, less obvious examples of features we consider sense-making data types on today’s web include flagging content as sensitive or inappropriate,

bookmarking, or muting a specific channel, account, tag, or type of content. All of these processes increase understanding of individual and collective preference, at the individual and collective levels, as these levels recursively feed into each other.

Today, these operations impact the visibility and usability of digital resources. What's more, they are subject to platforms' own profit-motive. Collective sense-making becomes subordinate to this motive and therefore also becomes more easily exploited by advertisers, aggressive state- and non-state actors, and platforms' own imperative to keep users hooked.

In the Neighborhoods ecosystem, these processes are decided by communities — neighbourhoods — and we aim for them to be able to vary widely from group to group, therefore we've abstracted these operations into a simple but powerful set of generic primitive components that can replicate what web2 offers today and much more!

In simple terms, community activators decide how specific sense-making data types impact the visibility of resources. Concretely they define if a context has access to a specific resource (defined through thresholds) and how that context displays that resource (defined by specific ordering algorithms). In our research neighbourhood example this can be seen in how only 'finished' articles have access to the public-facing context available for all members under the 'readers' role. Inside of this context, the best rated articles are shown first and take up more space in the UI. Other visibility operations include things like 'blurring' resources which are categorised as sensitive through reactions or hiding spam responses in a comment section.

Generally speaking, sense-making data types can be categorised as either qualitative ones such as **reactions** and **assessments** and quantitative ones such as measurable records of activity (e.g. number of posts and their associated timestamps), and the results of processes such as the winner of a game or the fulfilment of a promise. These latter are a part of applied **sensemaking methods**. Let's take a look at these!

Reactions

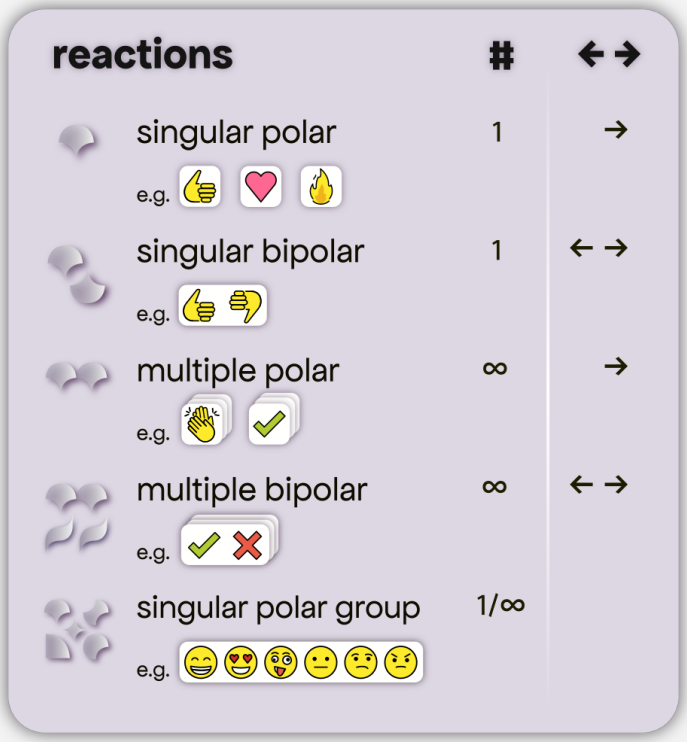

These are the most basic sensemaking data types and are equivalent to what we commonly find in today's platforms in the form of likes, hearts, upvotes/downvotes and other emoji-based reactions to resources. Through them agents can express an impression, assert an observation, signal appreciation or rejection among other diverse forms of acknowledgments.

Different neighbourhoods should be able to make cosmetic customizations of their sensemaking data types, because use of identifying symbols can impact a sense of in-group belonging and can boost participation. Maybe a surf community wants to call their 'like' equivalent 'waves' and use a sea-wave emoji.

While some differences are cosmetic, like the research journal neighbourhood using thumbs or arrows for gathering "yea/nay" expressions of preference for a proposed edition, some are computationally different. Reactions also have different degrees of complexity:

The purpose of setting these out as standards is to support cross-neighbourhood bridging. If two communities categorize their 'like' equivalent as such (even if they are cosmetically different), and use similar sense-making operations for them, it will be easy to bridge resources between two different neighbourhoods when so desired. We expect other forms of bridging to require more human mediation, such as bridging between two very different groups of emoji reactions.

MVP Detour

In our MVP demo, a very basic neighbourhood configured around memes as a resource and a linear feed as a context used two different multiple polar reactions: ‘lulz’ that allow members to signify ‘this meme makes me laugh’ and ‘mindblownz’ that signifies ‘this meme blows my mind’. Since these were implemented as multiple polar reactions, members could react to a single meme with 3 mindblownz and one lulz if the meme is mostly reflective but also a little bit funny! This way one can express a nuanced reaction to a specific meme aiding in the collective sensemaking around them. The feed context could be ordered through different algorithms, prioritizing either mindblownz or lulz, or creating specific combinations where one is valued at a fraction of the other.

Today, the difference between reactions like these and folksonomic, filtering categories known as *tags* is whether these are defined by, and already “baked into” a platform, or defined by any user whatsoever to support search and discoverability. The attachment of an additional layer of meaning to a resource is fundamentally the same, however. In neighbourhoods, the availability of a given additional layer of meaning is a question of group governance. For example, one could specify a threshold at which a tag is so useful for the group’s purpose that it becomes a more available reaction to be used across resources.

In our research journal neighbourhood example, tags help understand the topics of articles, and aid in the correct allocation of the 2X value that ‘topic specialists’ have when assessing articles that correlate to their area of expertise. If a neighbourhood decided to make a topic into a longer-term research track, tags could become staid reactions that readers use to help sort incoming submissions or articles from other journals. Alternately, members of any tier could sort their *personal* view of the neighbourhoood by a particular tag or reaction depending on their priorities.

Assessments

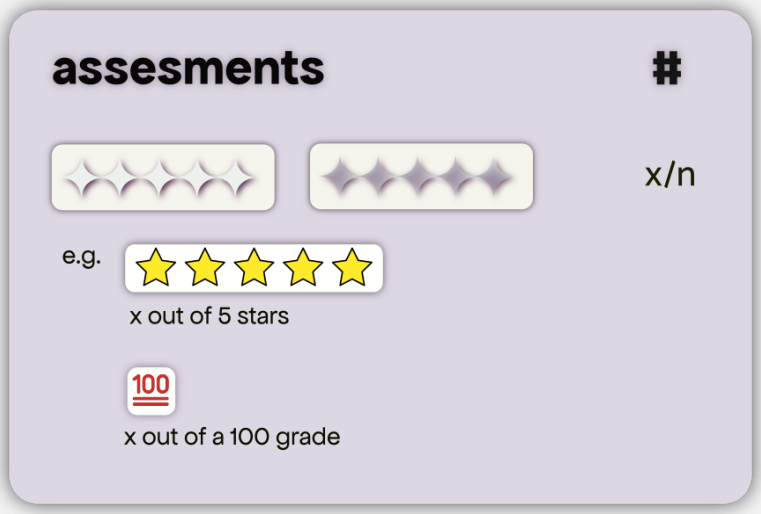

Assessments are a more sophisticated sense-making data type intended for more thoughtful forms of interaction between agents and resources. They are evaluative responses like ratings, grades and other structures in the form of x out of y. They are distinct from reactions in their capacity to be reacted to (in this sense, they become sort of a sense-making resource), as well as in having additional fields attached to them like a written comment that accompanies a given rating.

We find it important to distinguish between reactions and assessments, at a user experience level, in terms of how significant they feel to a community. While reaction can be seen as a casual, 'thinking fast' process, assessments should feel as a more thoughtful, 'thinking slow' one which is the reason why we envision them being often paired with a written comment. We also think it will be incredibly useful to have the capacity to react to assessments, as that can act as a further layer of collective sense-making.

For example, in our research journal neighbourhood example, assessments determine how the articles flow from 'draft' to 'candidate' to 'finished'. When 3 peer-reviewers have assessed a draft, and the average rating is larger than 4 stars, then it becomes a candidate. Subsequently, when a candidate is assessed by 3 edition editors and the average rating is larger than 4 stars, the article transitions to 'finished' and becomes publicly available. Even though only peer-reviewers can assess drafts and only editors can assess candidates, all members of the neighbourhood can react to those reviews, slightly affecting the final rating.

Sense-making methods

Sense-making data types would have only limited utility if not for the possibility of running various computations on them that can be fed back into the neighbourhood. This is all the more important given the multi-dimensional inputs and reactions neighbourhoods enable. In addition to summations, counts, and averages, we attend carefully to events and ranges which may trigger shifts in role or context that directly affect agency (as described above).

Creating and implementing sense-making methods is a huge design-space with ramifications into statistics, data analysis, governance and organisational practices. We consider this to be a vast topic which we are just beginning to scratch the surface of. It is here where neighbourhoods distinguish clearly from the algorithmic black-boxes of platforms, as the sense-making methods will not only be open-source but also developed by the community and implemented by each neighbourhood separately.

Scaling, Weighting & Limiting

Scaling is when a data type's base value is multiplied by a certain factor so that it is worth more relative to other data types. For example, in our research journal community, members who hold a topic specialist badge see the value of their assessments doubled when interacting with an article that corresponds to their area of expertise.

Limiting is, well...a limiting operation on sensemaking data types, modulating their possible uses. How many reactions can I perform in a given time? Or for a given set of posts? Such limitations could be confined to particular roles, helping distinguish them, or simply contribute to a neighbourhood's culture (say, capping the number of reactions from a specific account or implementing a decision as reversible within particular timeframe).

Statistics

One special, more general affordance of distributed socials is the possibility of enabling “opt in” analyses of samples of the neighbourhood. When Facebook ran its “emotional manipulation” study it was a feat of social science, albeit it done illegally and unethically. Neighbourhoods have the option of running any number of comparisons, or testing out various hypotheses, using sampling on categorical or quantitative data at the neighbourhood level or even at the agent-level.

Network-wide analysis is possible via *analysis of resources in shared contexts.* While formal voting may be strongly suited to some situations (and as popular as it has become in the web3 world), analyzing *object relations*, particularly the links between agents and resources, provides a day-to-day view of preferences, choices, and real activity that can inform group-wide decision-making about how to improve or manage conflict in a neighbourhood. In this sense, neighnourhoods can enable proto-governance, where decisions emerge and are sourced through agents interacting with each other and with resources.

In essence, the sky’s the limit for things like testing hypotheses, comparing activity to theoretical constructs, inferring relationships, and discovering associations between members, events, and actions. A neighbourhood may, for example, want to know if there was any association between a certain type of preference or activity and various tiers of membership or badge holding. They may want to understand the frequency of use of a certain module compared to another similar one to determine a group preference that can become a standard. One could compare groups, compare choices and preferences, or the averages thereof by using relevant, anonymized personal data. A neighbourhood focused on member achievement (a skill share group or an accountability group) also has the option of comparing a members’ activity to group norms or to their past actions.

Coming up next, we'll briefly sketch our approach to ***portability*** of social data, as well as ***bridging*** between communities. In a sense, bridging *is* portability at scale...

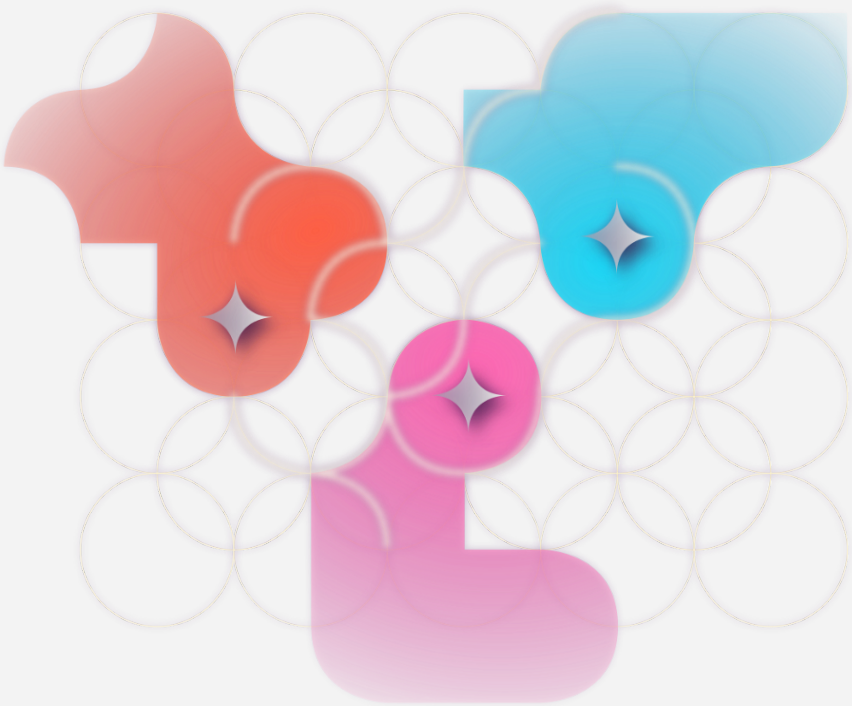

Portability

Portability refers to the way in which an agent's credentials (which reflects their standing in the neighbourhoods they are part of) are viewed differently in different neighbourhoods, and the rules each community uses to re-contextualize this data.

To fulfill on **portability**, it is important to have conventions on sensemaking data types and context-resource-scoring patterns that community activators can choose to adhere to. This will simplify the interoperability within and across different neighbourhoods. Using rep_lang, our Turing-complete programming language for social data, it is possible

to create simple standards for common data types like ‘hearts’, ‘upvote / downvote pair’, ‘claps’ and other simple reactions and scoring methods.

A note on the contentious notion of “reputation” taken across context: We do not think of membership, roles, badges, awards or achievements as ‘reputation’ in themselves. Instead, we see these as the specific social-data abstractions which communities use to formalize their culture to promote certain behavior, grant rights and share responsibilities.

In our view, reputation is the contingent social phenomenon that emerges from the interpretation and implementation of credentials. For this reason, we understand credentials as community-bound.

Therefore, it is always up to other communities to decide if and how they want to honor and translate other neighbourhood’s credentials. Additionally, it is ultimately up to an agent to decide if they want to reveal credentials from other neighbourhoods when entering a new neighbourhood. A contextual approach to reputation is Neighbourhoods response to the black mirror nightmares of universal reputation systems that collapse these complex social interactions into one-dimensional yet universal metrics.

By expanding the creation of nuanced and multi-layered data that changes in meaning within each community, we acknowledge the complex web of trust, reputation and culture unique to each community. Using our research journal neighbourhood as an example, members coming from reputable niche-research communities could immediately acquire a ‘topic expert’ badge when joining the research journal neighbourhood, without having to ‘build’ their reputation from scratch. This recognition of external credentials reflects a degree of trust placed by members of one neighbourhood in the reputational integrity of another neighbourhood. We call this process bridging.

Bridging

Bridging allows for a new kind of coordination and information sharing, without having to conform to a general overarching consensus. It refers to the referencing, or sometimes the importing of data from neighbourhood to neighbourhood. The possibility of bridging requires us to think about how communities share data and for which purposes. In the context of bridging, neighbourhoods implicitly bestow legitimacy upon other neighbourhoods' rules for generating social data (or take them with a grain of salt, referencing but not importing them, or trusting them but to a lesser degree than the nh's own modus operandi).

For example, a neighbourhood that aggregates articles only about a specific topic, like case studies on blockchains used in banana supply chains, could pull this niche content from more general, distributed ledger focused research journals. In return, they might grant the author(s) a portable "author badge" that reflects their contribution to the new context. Depending on the permissions used in the first neighbourhood, this whole process might require explicit consent of the author(s) — quite unlike the grimey web scraping of Academia.edu or Researchgate that gives nothing to authors except a redundant profile page to maintain that they never wanted in the first place!

But bridging also describes phenomena like the 'merging' or 'forking' of communities. We envision neighbourhoods, who find themselves having a high affinity, to be able to merge their communities into a larger group. As well as different ways in which the forking of a neighbourhood can create a parallel smaller community that is able to take its social-data with it at the time of the splitting. With a bridging pattern, we're effectively moving towards a 'memetic-mesh' network that generates long-reaching agent-centric feedback loops, beyond the boundaries of individual communities. In doing so, we enable large-scale coordination, without neighbourhoods having to, become, 'big' in order to do broader sensemaking.

Communities are therefore able to maintain their specific culture, norms and nuance, while still reaching out to other communities. All bridges exist **between** communities. There is no universal ledger where data is stored and ported from. While we do not yet have implementation details worked out, we embark on a process of tinkering and experimentation that will lead to stable & usable patterns for coordination across neighbourhoods.

Conclusion: Infrastructure for Imagination

Through our example “research journal neighbourhood” we are outlining a system for permissions based on members and roles, identification of different types of reactions, preservation of ‘resources’ presented in various ‘contexts’, subject to layouts and ordering algorithms chosen by the community of use for the fulfillment of its purpose(s).

This system is surprisingly simple. Yet, we intend it as a powerful framework for grasping how social networks and platforms are structured today and how some of that functionality can be replicated in a modular way with the tools Neighbourhoods is building. More than anything, the ontology presented throughout provides language for imagining your own neighbourhoods implementations.

In these early days, our ontology for social primitives and sensemaking data is, in many ways, infrastructure for imagination.

While we are developing with these concepts, the best way to put the ontology to use at present is to bring your own needs. We invite you to join us in ideating on prospects and potentials for future neighbourhoods through your curiosity and your desire.

We deeply appreciate your readership. If any of this inspired you, come join us on Discord, follow us on Twitter or go get yourself some $NHT ❤