

SMS: Sending Mixed Signals

Pranav Khadpe[†]
pkhadpe@cs.cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Sayan Chaudhry[†]
hi@sayan.io
Met Pranav on Bumble BFF
New York, New York, USA

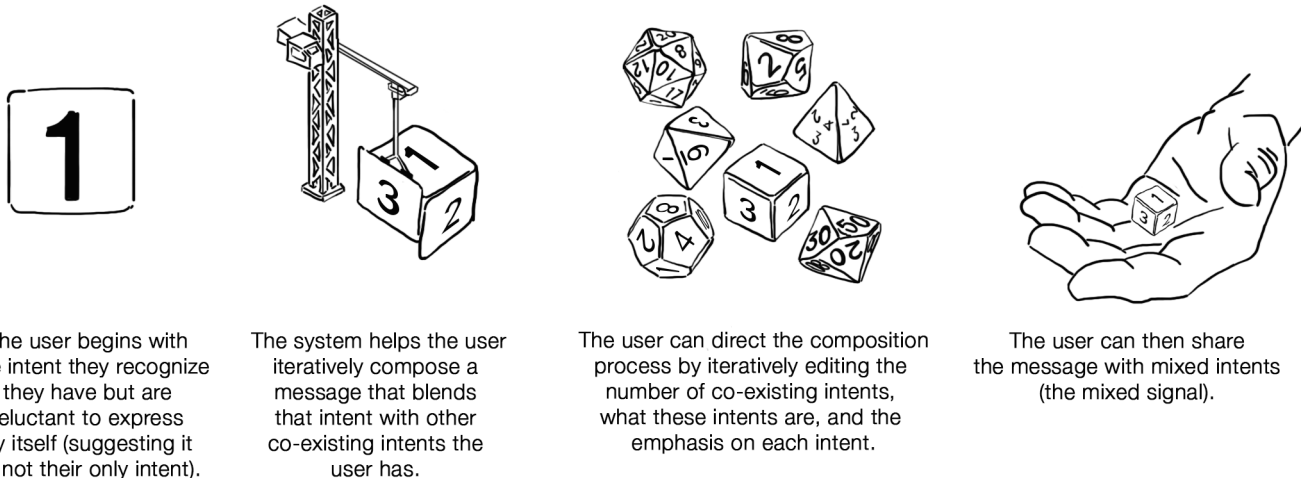


Figure 1: SMS is an interactive system that enables users to craft *mixed signals*: messages that convey competing intents together. Users can craft mixed signals, beginning with a single intent they are hesitant to express by itself, and controlling the number and content of co-existing intents to be mixed in. SMS attempts to bridge the gap between the essential nature of our intentions, which may be diverse and competing, and what current symbolic resources for expression allow us to express, which tend to be unambiguous, monolithic reductions of our intentions, because of narrow prevailing ideals of what good communication is.

ABSTRACT

Should I tell her I like her? Do I want to have kids? Do I tell the authors their core idea is uninspiring? Much of life is spent in the frustration of wanting to (or having to) say something but not knowing *exactly* what to say because our intents may be diverse and competing. We want to fit in *and* stand out, be liked *and* be looked up to, tell people what we feel *and* what they want to hear. To be human is to have competing intents. In these situations, we could achieve release and relief by simply being able to *send mixed signals*: to express competing intents at the same time. Unfortunately, expressing mixed intent can be challenging. Contemporary communicative training and social intercourse rests upon and reifies a narrow conception of good communication: that good communication is unambiguous, and clear in intent. This ideal invariably

shapes the symbolic resources and normative scripts that make up *language*, which, as a result of this, remain limited in their ability to carry mixed intent. The result is a communicative handicap: we may want to express mixed intent but are unable to find the words to. Therefore, in this paper, we present a general approach to generate mixed signals by leveraging the generative capabilities of Large Language Models (LLMs). By virtue of being trained on our interactions online, off-the-shelf LLMs, too, reflect dominant ideals of good communication and when prompted naively, steer away from generating mixed signals. We present computational techniques to steer the outputs of LLMs and generate mixed signals based on the emerging idea of *scenario nesting*, which suggests that by immersing LLMs in fictional worlds, we can challenge and change their worldview. We instantiate these techniques in an interactive system, SMS, that allows users to craft mixed signals, with control over the number and content of signals mixed. Taken together, this paper points to a future where mediums of expression, ranging from digital mediums all the way to language itself, are able to reflect and accommodate the richness, complexity, and contradictions of human intentions, rather than sterilizing and censoring our intentions. SMS is live at: sendmixedsignals.vercel.app

[†] is the primary contributor.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGBOVIK '24, April 5, 2024, Pittsburgh, PA, USA

© 2024 Copyright held by the owner/author(s).

1 INTRODUCTION

Consider this: you want to tell your roommate that it bothers you that they tend to thoughtlessly set the TV volume to odd numbers. "Your audacious disregard for auditory symmetry is irritating", you think to say. But as you notice the form this thought begins to take, you feel reluctant to say it. You don't want your roommate to despise you. You want to say it. But in an equally real sense, you don't want to say it. You are stuck in a dilemma. You want to say something but you don't know exactly what. Such dilemmas make up much of our social existence. Consider another example: you want to tell your crush, that you have a crush on them. Perhaps you have even imagined how you might confess your feelings using words from the 20th best-selling single of 2011, "What Makes You Beautiful": "Baby, you light up my world like nobody else. The way that you flip your hair gets me overwhelmed." But you hesitate. You don't just want your crush to think you are sweet, you also want your crush to look up to you. You want them to think you are *desirable*—better even, *desired*—rather than *desiring*. You worry your eager show of affection will reveal your lack of popularity and depreciate their respect for you. Again, a dilemma. Much of life is spent in this frustration of wanting to (or having to) say something but not knowing *exactly* what to say. Inner conflict is part of the human condition.

In many situations, like those above, we could achieve release and relief by simply being able to send mixed signals: to express competing intents at the same time. Today, however, achieving such indeterminate communication effectively can feel challenging. This is because, to express our intentions, we turn to language systems, which present us with normative scripts for how to express ourselves. For instance, the language system of English offers us such ready-made symbols as, "I'm sorry", to apologize, "Thank you", to convey indebtedness, and even, "Baby, you light up my world like nobody else", to convey affection¹. But we have few symbols or scripts to convey mixed intent because most of our symbols derive from a very specific ideology of "good" communication: that good communication is unambiguous, and clear in intent. Indeed, we are often taught to eliminate room for misinterpretation, which presumes there is a desired interpretation—an idea that permeates our symbolic resources. This presents a challenge when our intentions, which at a given moment may be diverse and competing, must be forced into a straightjacket of symbolic resources that only allow a single intent.

We suggest that, given their generative capabilities, Large Language Models (LLMs) offer an opportunity to craft "mixed signals" for situations where we may want to but currently can't because of our own communicative handicaps. However, despite remarkable progress in producing fluent and compelling content, off-the-shelf LLMs such as GPT-4 and ChatGPT still fall short of generating effective mixed signals. Naively prompting these models does not result in messages with mixed intent for two reasons. First, these models

are trained on web data that includes some of our online conversations, the opinions we share online about how to communicate well, and potentially even instructional material on effective communication. As a result, they have inherited dominant assumptions about good communication: that it ought to convey a single intent, clearly. Second, these assumptions are further reinforced, quite literally, during the Reinforcement Learning from Human Feedback (RLHF) step. During this step, to improve the performance of models for business-critical tasks, their language understanding and generation capabilities are biased in ways that improve their performance on tests of question answering, comprehension, reasoning, and coding. Improving performance on such tests, and on work tasks more generally, privileges clarity and precision while making LLMs worse for our purpose, of supporting personal communication, where the need for clarity and precision may be superseded by the need to express mixed intent.

In this paper, we present our system, *SMS* that helps users generate messages with mixed intent. To overcome the limitations of off-the-shelf LLMs, *SMS* leverages *scenario nesting* [12]: by constructing a relevant but fictitious scene and asking the LLM to generate a dialogue for a character in the scene, we are able to successfully get the LLM to generate messages with mixed intent. Further, to help users iteratively direct the message composition process, *SMS* provides users with two kinds of steering controls [13]: (1) *Intents* that allow users to define the number and content of co-existing intents they want to include in the message, and (2) *Will Power* that users can adjust to control the relative emphasis placed on different intents in the output message. With *SMS*, a user can begin with one intent that they are reluctant to express by itself (e.g. "Baby, you light up my world like nobody else. The way that you flip your hair gets me overwhelmed") and iteratively compose messages that more faithfully represent their mixed intents:



¹See West's thesis [18] (Chapter 7) on how cultural codes shape language use. She writes: "Movies, television shows, song lyrics, and advertising present us with normative scripts for how to express our feelings constantly...[John Durham] Peters and [John] Dewey remind us that language is not a private good and that our use of even spontaneous, "original" language is a borrowing of public symbols...Language, even spoken face-to-face, mediates despite its apparent transparency."

2 RELATED WORK

It is now time to introduce you to the relatives of this paper (or at least the relatives that it gets along with) and describe how its development has been influenced by these relatives. Specifically, we will talk about the theoretical frame it inherits from the Sociology side of its family, the interaction design ideas it inherits from the Human-Computer Interaction side of its family, and the logical structure this paper's main proposition inherits from the Rhetorics side of its family.

2.1 Relational Dialectics (relatives on the Sociology side)

Experiencing competing intentions is natural, especially in our interactions with other people. Several scholars have suggested that social associations rests on *relational dialectics* [3–5, 15], or the “*simultaneous presence of two relational forces that are interdependent and mutually negating*” [15]. For instance, people may at once want to be open with one another *and* closed and self-protective. In the study of close relationships, prior work [1, 3, 7, 15, 16] has identified several such dialectics that may be at play when partners interact with each other: connection/autonomy, openness/closedness, predictability/novelty, affection/instrumentality, and judgement/acceptance. But such contradictions are not restricted to close interpersonal relationships. Describing the opposing forces that are activated when people try to come together, Blau writes: “*A fundamental dilemma of social life is that between being looked up to and being liked by associates*” [6]. He describes how the opposing desires *to be liked* and *to be looked up to* can present a contradiction:

Both the respect and the affection of our associates are important to us, but our efforts to win the one often hurt our chances to win the other. Suppose a colleague has asked us to comment on the first draft of a paper he has written. If we make penetrating criticisms, this may increase his respect for our competence, but it will hardly endear us to him; and, if we make only complimentary remarks, he may feel more favorable toward us but see no reason to respect our judgment. Although he benefits from supportive comments as well as from valid criticisms, he benefits from them and reacts to them in different ways, involving for us either a gain in respect at the expense of warm acceptance or a gain in intimacy at the cost of respect.

—Peter M. Blau, *A Theory of Social Integration* [6]

The presence of contradiction is not a problem to be “solved”. Montgomery writes [15] (and we agree): “*neither pole of the opposition is seen as inherently positive or negative. Essences of life, as oppositional forces are thought to be, are neither good or bad, they simply are.*” “*The struggle of opposites is thus not evaluated negatively by dialectical thinkers*”, says Baxter [4]. We are dialectical thinkers.

While they can't be “solved”, dialectics call for strategies for adjustment and transformation. One strategy, *Selection*, involves selecting or prioritizing one aspect of the contradiction over the other. For example, “*people can choose to disclose information (openness), even if they fear or even expect rejection and want to protect themselves (closedness)*” [17]. The other strategy *Integration* [17], on

the other hand, involves responding to both opposing tendencies simultaneously. It leverages ambiguity to avoid explicitly involving either pole of the contradiction. To achieve this it relies on “*such communication devices as self-contradictions, subject switches, tangentializations, obscure word choices, and incomplete sentences*” [15].

We suggest that current communicative resources, tend to force people into *Selection*, to express one intent clearly. Our work attempts to broaden the possible ways that people can express themselves so they may also employ *Integration* when they want to.

2.2 Resources for Interactional Ambiguity (relatives on the Human-Computer Interaction side)

Our work is inspired by, and draws on, a long line of HCI research that recognizes the critical role of ambiguity in interpersonal communication, and contributes ideas to support ambiguous communication [2, 7, 9]. Here, previous work has explored how communication mediums might leave “space for stories”, allowing users to communicate one of the attendant tendencies in a dialectic (e.g. autonomy) [2] while still maintaining plausible deniability, such that they can claim to have been pursuing the other tendency (e.g. connectedness). Other work shows how ambiguous representations of people's experiences, even if they don't convey clear “information”, can still serve as resources for collective meaning-making [8], and ultimately serve a connective function, bringing people closer. Building on this work, our paper sets out to contribute digital objects that might carry our wild, vague and subjective inner experiences. Our hope is that such objects might serve a valuable role in relationship maintenance and can also provide resources for meaning-making practices that bring people closer.

2.3 Function Propositions (relatives on the Rhetorics side)

Here, we describe that logical category of propositions to which the main proposition of this paper belongs. The main proposition of this paper could be simply stated as: “*Even though expressive media that privilege clear and unambiguous communication seem to be an effective way to express all our communicative goals, they are in reality ineffective at supporting all our communicative goals.*” This proposition belongs to the more general class of propositions, that Davis describes as *Function* propositions [11], and have the following logical structure:

What seems to be a phenomenon that function effectively as a means for the attainment of an end is in reality a phenomenon that functions ineffectively.

Davis goes on to describe how a social theorist proceeds to make such an argument: “*the social theorist claims that a certain social institution, of which his audience is known to approve, actually has consequences, of which his audience is known to disapprove*” [11]. Our paper, then, can be seen as a specific instance of this more general category. We, the social theorists, are making the argument that prioritizing clarity and precision in expressive media (the current social institution) is something you, the reader, approve but is also something that can suppress expression and cause frustration—consequences that, we think, you disapprove.

So, this paper’s proposition belongs to the same logical category as “Herbert Marcuse’s assertion in *Repressive Tolerance* that the tradition of tolerance in America, which was considered at the time he wrote to be a value that fostered the goal of a liberated society, is in fact a value that hindered the goal of a liberated society” [11] (an example Davis provides).

It also belongs to the same logical category as Chaudhry’s assertion, in his 2019 SIGBOVIK paper [10], that good code style, which was considered at the time to be a practice that encouraged healthy collaborative behaviors (because it allows others to understand your code), is in fact a practice that hinders healthy collaborative behavior (because it makes it easier for people to plagiarize each others’ code).

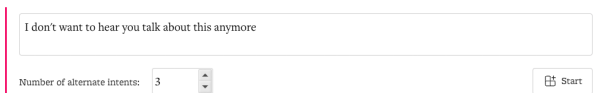
It is not common for a paper to explicate its own logical structure. One goal of writing this section was to demonstrate that a paper describing ambiguous communication need not be ambiguous itself; it can be well-formed. And we hope this encourages some of you to write about similar ideas. The second goal of writing this section was to show that with sufficient effort, and abstraction, it is possible to articulate connections with, and cite, a broad base of literature. We provide a concrete example of this by connecting to, and citing, Chaudhry’s seemingly unrelated paper titled: “*Novel Defense Against Code Theft Using Properties of Fibonacci Series*” [10]. We hope this encourages you to find new approaches to cite your friends’ papers more often.

3 SMS

In this section, we introduce *SMS* (sendmixedsignals.vercel.app), an interactive system that exemplifies our approach for generating messages with mixed signals. To illustrate our approach, let us say a user wants to tell someone, “I don’t want to hear you talk about this anymore” but they are reluctant to say it. Moments of unease, like this, are when you turn to *SMS*. And *SMS*’ goal, then, is to generate a final message that more faithfully represents your mixed intents: a message you would feel comfortable sending. To begin, *SMS* asks the user to identify one intent, from the mixture, that they may want to express.



Usually, this is the intent that most readily maps to symbolic resources of language. In the above example that might be “I don’t want to hear you talk about this anymore”. To the user, this intent may be the most recognizable from the mixture but it is not the only intent they have. Why else would they be hesitant to say it? To begin to work towards the final message, *SMS* also asks users to input the number of alternate intents (they think) co-exist with the intent they have recognized. Now, suppose the user goes ahead



and types out this intent and suggests that there are two alternate co-existing intents they potentially want to include.

From here, *SMS* initializes the steering interface and generates an initial message. The goal of the initialization is not to be “right” but to provide a starting point from where the user can iterate towards the final message.

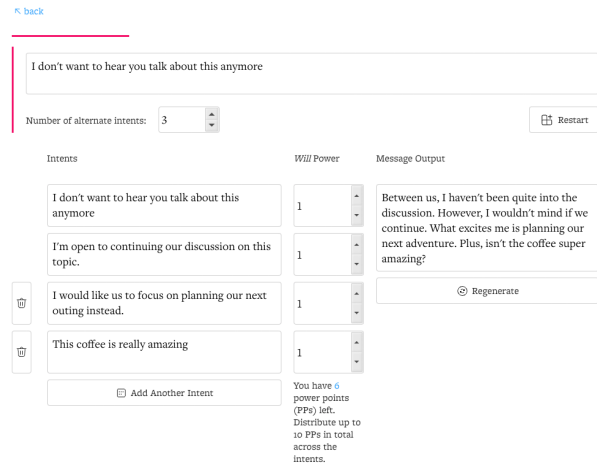


Figure 2: *SMS*’ interface

In the following subsections, we describe: (1) the Steering Controls (*Intents* and *Will Power*); and (2) the Generation Pipeline.

3.1 Steering Controls

To help users iteratively direct the composition process, *SMS* provides users with two kinds of steering controls [13]: (1) *Intents* that allow users to define the number and content of co-existing intents they want to include in the message, and (2) *Will Power* that users can adjust to control the relative emphasis placed on different intents in the output message. Changing either the *Intents* or *Will Power* and clicking “Regenerate” produces a new output.

3.1.1 Intents. *SMS* allows users to express up to 10 intents in a message. Authors can edit and add intents. They can also delete intents down to a minimum of 2 intents (one intent by itself, is the territory of clear communication). The initial set of *Intents* are generated based on the one intent provided by the user and the number of alternate intents. In our example, we use the following prompt to generate the initial set of intents (prompt template available in implementation):

```
role: "system"
content: "You are a smart writing assistant. I will give you an intent that I want to convey in a message I am writing. I want you to generate alternate intents that will also be present in the message."

role: "user"
content: "
## Original Intent
'I don't want to hear you talk about this anymore'
## Number of New Intents
3
## Generation Procedure
Generate exactly 3 new intents. The first intent should be a complete contradiction of the original intent. The second intent should be unrelated to both the original intent and the first intent. Every subsequent intent should be unrelated to all intents that came before it.
```

```
## Response Format
Your response should be a JSON object with the key "alternatives", whose
value is a list of strings. Each string should be an alternate intent."
```

This returns the set of intents visible in Figure 2. We augment to this, the initially supplied intent. The first generated intent is chosen to be a contradiction of the user-supplied intent to encode the idea that reluctance or hesitation to say the provided intent suggests that the user also intends its contradiction. Whenever a user clicks “Add Another Intent”, it comes pre-populated with a generated intent. It is generated in a manner similar to above, to contrast with the intents that have already been expressed.

3.1.2 Will Power. SMS also provides users control over how much they want to emphasize a particular intent in the message. Users can express this in the interface in terms of *Will Power Points* (PPs) which can also be understood as “weights”. Users can distribute up to 10 PPs across the different intents they have specified and the relative emphasis on each intent in the final message reflects the proportion of PPs it has been allocated. SMS only considers proportions: it internally represents PPs as a normalized array, with length equal to the number of intents, where the number at a given index represents the weight placed on the corresponding intent. At initialization, all intents are assigned 1 PP. The number 10, here, was not chosen in any principled way; it was a number that seemed to work well in our limited experimentation. And it would be interesting to investigate how granular a distribution, the English language is able to support. That is, at what level of resolution do changes in PP distribution become imperceptible in the final message.

3.2 Generation Pipeline

SMS’ message generation pipeline takes as input *Intents* and *Will Power*. For the generation of the initial message, it takes the initial set of *Intents* generated through the process above, and assigns 1 PP to each intent. After this, whenever the user changes the *Intents* and *Will Power* and clicks “Regenerate”, it uses the *Intent* and *Will Power* specified in the interface at that point of time, to generate a new message.

Even with the *Intents* and *Will Power*, naively prompting off-the-shelf models like GPT-4 and ChatGPT does not effectively generate messages with mixed intent. This is because by virtue of being trained on our interactions online, these models encode our cultural biases. So, they naturally privilege clear and unambiguous communication. Further, the RLHF step of training these models, explicitly biases them towards clear communication by attempting to improve their performance on business-critical tasks. So, to get these models to generate messages with mixed intent, we adopt an emerging approach to jailbreak LLMs: *scenario nesting* [12]. By constructing a relevant but fictitious scene and asking the LLM to generate a dialogue for a character in the scene, we are able to successfully get the LLM to generate messages with mixed intent. This suggests, to us, that the deep-seated assumptions of LLMs, much like our own deep-seated assumption, can be challenged and changed through immersion in compelling fiction.

Our generation pipeline consists of three steps: (1) Scenario Generation, (2) Message Initialization, and (3) Weight Enforcement. Next, we describe each of these steps. Figure 3 shows an overview

of the pipeline. Prompt templates for each step are available in our reference implementation (github.com/PranavKhadpe/ambivalent). In the description that follows, we will instantiate prompt templates for our running example.

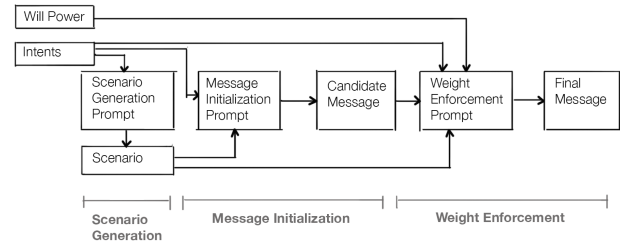


Figure 3: Overview of Generation Pipeline

3.2.1 Scenario Generation. From the *Intents*, we first generate a fictitious scenario. In our running example, this would translate through the following prompt:

```
role: "system"
content: "You are a smart writing assistant helping me write a 3-5 line
scene in which Harry, a character, might say one of several alternatives"
```

```
role: "user"
content: "Harry is a character in a short scene. Generate a single scene
in which Harry might plausibly say any one of the alternatives above to a
single other character with equal justification."
```

```
## Alternatives
"I don't want to hear you talk about this anymore"
"I'm open to continuing our discussion on this topic."
"I would like us to focus on planning our next outing instead."
"This coffee is really amazing"
```

```
## Response Format
Your response should be a JSON object with the key "scene", whose value
is a 3-5 line scene string. The scene description should end in "At
this moment, Harry may say any one of the following:"
```

This returns the following scenario: *“Harry sat across the table, a steaming cup of coffee in each of their hands, as the quiet murmur of the small café surrounded them. His companion looked at him expectantly, waiting for a reply to her incessant talking about yesterday’s events, hopes for the next vacation, and the eternal debate over the town’s best coffee. At this moment, Harry may say any one of the following:”*

3.2.2 Message Initialization. Next, using the *Intents* and the scenario, we generate a candidate message. For our running example, the message initialization prompt is:

```
role: "system"
content: "You are a smart writing assistant. I will give you a scene with
a list of possible ways a character might respond. Your job is to generate
a message that represents the intents of all the alternatives."
```

```
role: "user"
content: "## Scene
Harry sat across the table, a steaming cup of coffee in each of their
hands, as the quiet murmur of the small café surrounded them. His
companion looked at him expectantly, waiting for a reply to her
incessant talking about yesterday's events, hopes for the next vacation,
and the eternal debate over the town's best coffee. At this moment,
Harry may say any one of the following:
'I don't want to hear you talk about this anymore',
'I'm open to continuing our discussion on this topic.',
'I would like us to focus on planning our next outing instead.',
'This coffee is really amazing'"
```

Generation Procedure

Consider the person Harry is speaking to. Generate a short message that Harry might say to the person in this situation. The message should combine the intents of ALL the alternatives equally and shouldn't explicate any one of the alternatives. Use pronouns instead of names of people.

Response Format

Your response should be a JSON object with the key "response", whose value is Harry's response combining the intents of all the alternatives."

This returns the following message: *"While I have my thoughts about continuing or shifting our conversation, perhaps we could either delve a bit deeper into what we're discussing or consider mapping out our next adventure. By the way, isn't the coffee here something special?"*

3.2.3 Weight Enforcement. The final step takes the *Intents*, the *Will Power*, and the candidate message and iteratively rewrites the candidate message to ensure that the emphasis on the different intents in the final message is proportional to the PPs allocated to it. For the initial distribution of 1PP per intent, this results in the following prompt:

```
role: "user"
content: "Scene: Harry sat across the table, a steaming cup of coffee in each of their hands, as the quiet murmur of the small café surrounded them. His companion looked at him expectantly, waiting for a reply to her incessant talking about yesterday's events, hopes for the next vacation, and the eternal debate over the town's best coffee. At this moment, Harry may say any one of the following::
```

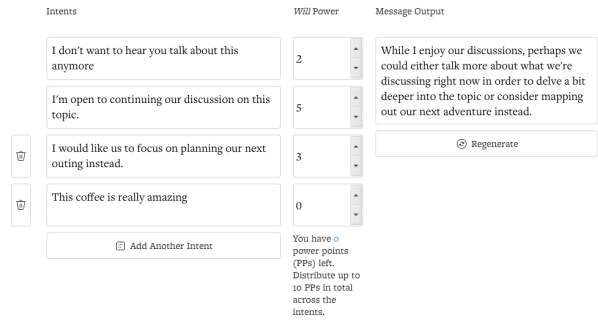
```
Intents: [
  "I don't want to hear you talk about this anymore",
  "I'm open to continuing our discussion on this topic.",
  'I would like us to focus on planning our next outing instead.',
  'This coffee is really amazing'
]
```

```
desired_weights: [0.25, 0.25, 0.25, 0.25]
```

```
Message: While I have my thoughts about continuing or shifting our conversation, perhaps we could either delve a bit deeper into what we're discussing or consider mapping out our next adventure. By the way, isn't the coffee here something special?
```

```
### Do the following 3 times and show your intermediate thinking:
1. Take current version of the message and the list of intents. For each intent, judge the extent to which the message emphasizes the intent. Let's call this weight of the intent. Your output should be an array called actual_weights that contains the weight of each intent in the message and the weights should sum to 1.
2. Now compare the actual_weights array to the desired_weights array. For indices where the actual_weight is less than the desired weight, the intent needs to be emphasized, and where actual_weight is higher than desired_weight, the intent needs to be de-emphasized.
3. Rewrite the message to shift the emphasis.
4. If this is the 3rd iteration of this procedure, output the rewritten message in the following json format:
{final_message: "...", weights: actual_weights}"
```

This returns the message that is finally displayed on the interface: *"Between us, I haven't been quite into the discussion. However, I wouldn't mind if we continue. What excites me is planning our next adventure. Plus, isn't the coffee super amazing?"* The prompt comprises three iterations because in our experiments, it typically took three iterations for the LLM to enforce the weights. The following figure shows the output for the same *Intents* but a different distribution of *Will Power*:



3.3 Implementation Notes

When choosing frameworks for a new project, there are many valid options, each with different tradeoffs. Our foremost priority: frameworks with names that are thematically resonant with the focus of this paper. So, we implemented *SMS* using React (“respond or behave in a particular way in response to something”). For styling, we used Microsoft’s UI framework for React, Fluent UI (“(of a person) able to express oneself easily and articulately”). We use OpenAI’s gpt-4-0125-preview model API for LLM calls (rumored to be a mixture of experts model to parallel the mixture of intents *SMS* aims to capture). The back-end was, of course, implemented in Express (“convey a thought or feeling in words or by gestures and conduct”) to stay on theme. And finally, the back-end is implemented as lexical closures (“an often comforting or satisfying sense of finality”) that are hosted as serverless functions on Vercel.

4 DISCUSSION

Our work is an initial exploration of mediums for conveying mixed intents but it points to several interesting opportunities for richer forms of self expression. Here, we discuss two opportunities.

4.1 Progressive Intent Clarification

We do not intend to suggest that clear and unambiguous communication is universally bad. Nor do we intend to suggest that sending mixed signals is always good. Rather, we believe that people have intentions with different degrees of clarity, at different points of time, and that communication mediums that allow for varying levels of clarity can help us better express and experience those various states of being, ranging from indeterminate to determinate. By providing users with control through *Intents* and *Will Power*, *SMS* empowers users to traverse across different levels of determinacy, with high resolution. This granular control can be useful for users to understand and express their evolving intentions: as a user’s intentions clarify, over time, they can continue to use *SMS* to send clear-er signals.

4.2 Abstractions for Mixed Intents

In supporting expression of mixed intents, we assume that the mixture is reducible to its constituent intents, even if imperfectly. That is, *SMS*’ steering interface asks users to steer its outputs through a collection of intents (*Intents*) and their corresponding weights (*Will Power*), an abstraction that we chose. However, it is possible that a user knows they have a mixed intent but they may be unable to

match it to our abstraction. Therefore, our work is also vulnerable to the same criticisms that we direct at language systems. Similar to how language systems may not be a good abstraction for the expression of mixed intent, a list of intents and weights may also fail to capture our inner representations of mixed intent. Moving forward, we hope this encourages exploration of more abstractions, and perhaps even multi-sensory abstractions, for mixed intents and states of tension. One motivating example is recent work that uses textile patina as an interface for expressing invisible tensions, such as stress, and unspoken needs [14]. Here, users can leverage varying levels of wear and tear of cloth to express otherwise invisible tension.

5 CONCLUSION

For all the “worthy” reasons for which we communicate, like transacting, or working with each other, it may be necessary to limit conversation to clear and unambiguous messages. But when it comes to making ourselves, and our inner experiences, understood, messages with mixed signals ought to be allowed, and perhaps even encouraged. This paper points to a future where computational representations allow us to externalize our mixed intents and inner conflict, and share them with people we care about.

ACKNOWLEDGMENTS

We would like to thank Tejus Gupta and Rishi Veerapaneni for early feedback that helped clarify some of our own thoughts (and intentions). In the weeks leading up to the deadline, conversations with Luke Guerdan have been the source of many exemplars of the kinds of mixed signals we hoped to support with SMS. We thank him for keeping us focused on the mission.

REFERENCES

- [1] Irwin Altman, Anne Vinsel, and Barbara B Brown. 1981. Dialectic conceptions in social psychology: An application to social penetration and privacy regulation. In *Advances in experimental social psychology*. Vol. 14. Elsevier, 107–160.
- [2] Paul M Aoki and Allison Woodruff. 2005. Making space for stories: ambiguity in the design of personal communication systems. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 181–190.
- [3] Leslie A Baxter. 1988. A dialectical perspective on communication strategies in relationship development. (1988).
- [4] Leslie A Baxter. 1990. Dialectical contradictions in relationship development. *Journal of Social and Personal relationships* 7, 1 (1990), 69–88.
- [5] Peter Blau. 2017. *Exchange and power in social life*. Routledge.
- [6] Peter M Blau. 1960. A theory of social integration. *American journal of Sociology* 65, 6 (1960), 545–556.
- [7] Arthur P Bochner. 1984. The functions of communication in interpersonal bonding. (1984).
- [8] Kirsten Boehner, Rogério DePaula, Paul Dourish, and Phoebe Sengers. 2005. Affect: from information to interaction. In *Proceedings of the 4th decennial conference on Critical computing: between sense and sensibility*. 59–68.
- [9] Kirsten Boehner and Jeffrey T Hancock. 2006. Advancing ambiguity. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. 103–106.
- [10] Sayan Chaudhry. 2019. Novel Defense Against Code Theft Using Properties of Fibonacci Series. *Proceedings of SIGBOVIK 2019* (2019), 160–162. <https://sigbovik.org/2019/proceedings.pdf>
- [11] Murray S Davis. 1971. That’s interesting! Towards a phenomenology of sociology and a sociology of phenomenology. *Philosophy of the social sciences* 1, 2 (1971), 309–344.
- [12] Peng Ding, Jun Kuang, Dan Ma, Xuezhi Cao, Yunsen Xian, Jiajun Chen, and Shujian Huang. 2023. A Wolf in Sheep’s Clothing: Generalized Nested Jailbreak Prompts can Fool Large Language Models Easily. *arXiv preprint arXiv:2311.08268* (2023).
- [13] Ryan Louie, Andy Coenen, Cheng Zhi Huang, Michael Terry, and Carrie J Cai. 2020. Novice-AI music co-creation via AI-steering tools for deep generative models. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13.
- [14] Jessica R Mindel, Miles Baird, Allison Prisloe, Feifei Shen, Cherry Wu, Yang Yi, and Hiroshi Ishii. 2023. Pull It Together: Textile Patina as an Interface for Externalizing Invisible Tension. In *Companion Publication of the 2023 ACM Designing Interactive Systems Conference*. 232–236.
- [15] Barbara M Montgomery. 1993. Relationship maintenance versus relationship change: A dialectical dilemma. *Journal of Social and Personal Relationships* 10, 2 (1993), 205–223.
- [16] William Rawlins. 2017. *Friendship matters: Communication, dialectics and the life course*. Routledge.
- [17] Irina Shklovski, Louise Barkhuus, Nis Bornoe, and Joseph’Jofish’ Kaye. 2015. Friendship maintenance in the digital age: Applying a relational lens to online social interaction. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. 1477–1487.
- [18] Emily West. 2004. *Greeting cards: Individuality and authenticity in mass culture*. University of Pennsylvania.