On the Linguistic Behavior of Keysmashes

Allison Park

Abstract

Keysmashes may be defined as strings of letters, punctuation, and numbers that do not represent words in any language, yet seem to carry social meaning when used in text-based conversations (e.g. “hsgkjsdnsdf”). This paper serves to explore the linguistic nature of keysmashes, and characterize their forms. Experiment 1 circulated a survey to investigate whether keysmashes behave according to linguistic criteria, asking participants to perform such tasks as judging whether keysmashes have meanings, identifying what those meanings are with and without context, sorting social groups based on their likelihood to keysmash, and performing acceptability judgements. This survey found keysmashes to conform to five criteria for language; most notably, keysmashes exhibit semanticity, standards of form, and arbitrariness. Experiment 2 investigated these standards of form, examining the well-formedness of keysmashes with a constraint-based analysis. Eight primary constraints were hypothesized to have an effect on keysmash well-formedness: length, keyset, punctuation, vowel, lexical, phonotactic, repetition, and gesture. Experiment 2 circulated a survey that asked participants to make acceptability judgements. Results were analyzed with particular attention to the existence and size of effect of the hypothesized constraints, utilizing linear models and aggregates of average acceptability ratings. The constraints Punctuation and Vowel are violated when instances of what they are named for exist in a keysmash. The constraints Length and Repetition are hypothesized to function in a sort of bell-curve: having too little or too much of the described quality incurs violations. The constraint Keyset is violated when characters outside of the “home row” of a standard QWERTY keyboard are included in a keysmash. Similarly, the constraints Lexical and Phonotactic are violated when entire words (or clusters that follow English phonotactics to the extent they may initially parse as words) are present in a keysmash. Finally, the constraint Gesture is violated when the keysmash is primarily made up of characters on only oe side of the QWERTY keyboard as would be typed by a touch typist. Upon analysis, the hypothesized constraints vowel, length, keyset, punctuation, lexical and phonotactic proved to have statistically significant effects on keysmash acceptability ratings.

1. Introduction

A keysmash is a string of letters (and sometimes punctuation and numbers) that do not represent words in any language (e.g. “hsgkjsdnsdf”). Although they are named after smashing hands on a keyboard to create 'nonsense', modern keysmashes do not represent random conversational noise: these strings seem to be used to express feelings or convey tone in a consistent and understandable manner in some Internet subcultures. These qualities, particularly when keysmashes' forms do not seem particularly meaningful, call into question whether keysmashes behave linguistically, or are simply entropy or emotive conversational 'noise'. This work aims to investigate two key aspects of keysmash according to linguistic principles: Whether they are linguistic and behave as other accepted linguistic forms of expression do, and what grammatical standards, if any, govern the form of keysmashes.

To evaluate their potential linguistic nature, a survey was distributed online in Experiment 1 to investigate whether keysmashes behave according to certain linguistic criteria: arbitrariness of keysmash form, the social meaning and associations of keysmashes, and acceptability effects in and out of specific contexts.

Following this, a second experiment circulated another survey to investigate said standards of form. This second experiment compared the well-formedness ratings of keysmashes to the textual form of the keysmash using a constraint-based analysis, in an effort to explore whether there are specific patterns or constraints of form which make keysmashes read as more or less acceptable to regular keysmashers.

2. Literature Review

Investigating the possibly linguistic nature of keysmashes begins with criteria for linguistic behavior: what defines language as language? Hockett's 1960 article *The Origin of Speech* includes thirteen criteria that describe spoken language. Aside from eight acoustically focused criteria, Hockett lists arbitrariness, semanticity, productivity, displacement, and traditional transmission as distinctive qualities of human language. He defines displacement as the ability to "talk about things that are remote in space or time (or both) from where the talking goes on", and defines traditional transmission as the involvement of others in teaching the use of and traditions associated with the language (versus innate knowledge, as in birdsong and other animal communication). As keysmashes are not immediately and obviously productive as would be expected of a fully-fledged language, displacement was set aside in favor of examining other criteria.

While one cannot express new and complete thoughts entirely in keysmashes, observed use of keysmashes in online communities suggests that keysmashes do have meanings beyond shrieking that might be represented by random keyboard inputs. Keysmashes themselves are not quite random keyboard inputs: there are user intuitions about well-formed keysmashes versus ill-formed ones, suggesting standards of form. However, analyses of keysmashes are largely absent in the linguistic literature. Gretchen McCulloch's *Because Internet* includes a few observations about keysmashes, including mentions of patterns in keysmashes themselves and how they affect social acceptability of various keysmash compositions (McCulloch pp. 6-7). At least one other document points out some patterned tendencies in keysmashes, but not nearly to the detailed extent required for an analysis of either linguistic behavior or keysmash well-formedness, as its focus was on linguistic feature spread in Twitter communities (Ygartua

2018). Other papers that mention keysmashes at all tend only to say they are examples of new Internet-based language conventions, and do not begin to make the argument that they behave in patterns and have characteristics usually attributed to language.

Without literature or prior studies as reference points, the survey for Experiment 1 was created to gather data on relevant linguistic criteria: semanticity, standards of form, arbitrariness, and traditional transmission. Experiment 2 utilized findings from Experiment 1 as well as observations from McCulloch to formulate its hypothesized constraints; McCulloch notes that keysmashes tend to begin with “a” or “asdf”, tend to consist of characters on the home row of a QWERTY American English keyboard, and rarely contain numbers (McCulloch pp. 6-7). These observations, along with others gathered from Experiment 1, formed the basis for Experiment 2’s hypothesized constraints: length, keyset, punctuation, vowel, lexical, phonotactic, repetition, and gesture.

3. Experiment 1

3.1. Experiment 1: Introduction

Experiment 1 set out to investigate whether keysmashes behave according to linguistic criteria, rather than representing the meaningless noise their seemingly-random forms imply. To this end a survey was circulated online through various social media and chat platforms. The survey was based on Hockett’s criteria for language, focusing on six relevant criteria as could be adapted for typed communication: arbitrariness, semanticity, productivity, and traditional transmission.

3.2. Experiment 1: Methods

Experiment 1’s survey was divided into six sections of questions. The first half of the survey gathered basic data on participants and their background knowledge, asked why and in what contexts keysmashes are frequently used, and asked which groups frequently use keysmashes. The second half of the survey included two types of tasks, first asking participants for acceptability judgements with and without context, then asking participants for the meanings of various keysmashes with very little context.

3.2.1. Definitions and background information

The first section concerned itself with definitions: it asked whether the respondent knew what a keysmash is, how often they keysmashed, whether their friends keysmashed often, and whether they identified as a keysmasher. These were single-answer, multiple-choice style questions. This section was designed to gather basic demographic information about the respondents so that patterns in responses might be better identified and analyzed.

A variety of contextless example keysmashes, all instances found in public chat group conversation history, were included after the demographics section to provide some standard of what the survey considers a keysmash. These seven keysmashes included both common letters and letters that are not often found in common keysmashes. The longest keysmash included was 16 letters long; the shortest keysmash included was 5 letters long. Three of the seven used all uppercase letters, while the other four used all lowercase letters. None of the example keysmashes contained any numbers or special characters.

3.2.2 Keysmash meanings and contexts

The second section asked why and when keysmashes are used, alongside basic questions about what keysmashes express. The majority of questions in this section were checkbox style, allowing respondents to select as many options as they felt were true. Aside from gathering data about a variety of uses for and meanings of keysmashes, this style of question was instrumental in gathering data about what keysmashes are *not* used for and do not mean. Each checkbox question in this section included options that were hypothesized to be unacceptable or at least unattested in common keysmash use. The question "Could someone use a keysmash when they feel…" included "angry" and "sad" with the other four emotional options (surprised, flustered, amused, affectionate) and a non-emotional option ('None of these are valid uses for a keysmash'). These three particular options were predicted to be chosen much less often than the other four emotional options. Allowing participants to select several options allowed one question to gather data about both acceptable and unacceptable uses for a keysmash at once.

Other hypotheses tested this way involved the conversational contexts of keysmashes and the general question of whether keysmashes have inherent meaning. The question 'If you keysmash, when do you keysmash? Check all that apply.' included an option that specified keysmash use in professional and business contexts among the other varieties of use that were hypothesized to actually occur. Options in response to 'Keysmashes are commonly used for... (Check all that apply.)' included 'representing an acronym' and 'representing an action', both of which were hypothesized to be unacceptable. The question concerning keysmash semanticity was specifically phrased to generalize out of the personally-based content of the previous question, using the word "commonly" to imply usage by the majority rather than the particular respondent.

All checkbox questions, except the question asking "Could someone use a keysmash when they feel... (Check all that apply.)", included free-response options. These data were not included in the following analysis, but were used to gather information from respondents about potentially overlooked meanings and uses the survey did not include for future work.

One question in this section asked whether respondents felt they could differentiate keysmashes used for different emotions from each other. This question used a single-answer, multiple-choice format to force respondents to choose the answer they felt best fit themselves. It was particularly important to the hypotheses concerning keysmash semanticity that a distinction be drawn between the ability to distinguish keysmashes *with* context or *without*.

3.2.3. Keysmashes and social groups

The third section asked about keysmashes in relation to groups of people, and included three questions. The first question asked directly whether people from different social groups keysmash differently, while the last question asked whether reading a keysmash gave the respondent an idea of what groups that keysmasher might be a part of. These questions were separated by a grid question about different groups' likeliness to use keysmashes.

The grid question presented nine different groups for rating on a scale from 1 (a member of this group is less likely to keysmash) to 5 (a member of this group is very likely to keysmash). The directions specified that 3 should be chosen if membership in a particular group does not influence likeliness to keysmash. Groups included 'LGBTQ+ people' and 'Non-LGBTQ+ people' because keysmashing is popularly associated with LBGTQ+ people and communities online (Wikipedia). 'People under 25' and 'People in fandoms' were also included to investigate other existing popular associations with keysmashes. (People in fandoms were defined in the grid question's instructions as follows: "People who participate in creative activity related to media such as books, shows, movies, or games. People in fandoms might create fanworks relating to the media, consume these fanworks (reading fiction, commenting on art, blogging about the media, etc), or socialize with others specifically to discuss the media.")

As with the checkbox-style questions in section 2, the grid question included groups that were hypothesized to be either unrelated to or unlikely to use keysmashes; groups hypothesized to be unrelated were 'People with a specific gender identity', 'People with a specific political alignment', and 'People of a specific faith'. I hypothesized that specific identities would be less indicative of keysmash use or likeliness than the broader categories above. The specific categories were described in the instructions for the grid question. The survey specified that "The options 'People with [category]', and the option 'People of a specific faith', intentionally do not describe specific groups". These options were intended to capture particular, but unspecified, categories that could influence one's likeliness to keysmash.

The group 'educators' was hypothesized to be unlikely to use keysmashes. The group 'students' was added to contrast with 'educators', for a total of nine groups.

3.2.4. Acceptability judgements

The fourth and fifth sections presented keysmashes for acceptability judgements. The fourth section included ten contextless keysmashes, while the fifth section included seven keysmashes within a conversational context. Respondents were asked to rate how "well-formed, acceptable, natural, or like a 'proper' keysmash" they found the given items on a scale of 1 to 5, where 5 was the most acceptable and 1 was the least. Respondents were specifically instructed to choose 3 if they felt they could not judge the given keysmash.

Half of the contextless items were collected from actual utterances in public chat groups, while the other half were items created to investigate the acceptability of more unusual forms. The unusual items included a keysmash that utilized the special characters @, #, and $; a keysmash with characters from "edges" of the keyboard such as W, Q, O, I, and P; a very short (three-character) keysmash; a very long keysmash that included numbers; and a keysmash that included the vowels E and O near its beginning. These characteristics were hypothesized to be less acceptable than keysmashes composed of letters from the home row of a QWERTY keyboard.

The items with context were presented in a format meant to represent messages sent by chat users, as this format would be familiar to the online contexts where keysmashes are used in conversation. Each context involved up to two interlocutors, indicated with [User 1A] or [User 2B]. These user indicators were included on the same line as a message, so an exchange would look like the following:

> [User 1A] I'm so boooooored
> [User 2B] Hi so boooooored I'm [@User 2B]
> [User 1A] hdsfgjbb stop

Since the user tags look visually similar, the directions for the acceptability judgements with context included directions to note what user was sending a message, and specifically noted that some examples involved only one user sending several messages while others involved multiple users responding to each other. To avoid interpretations of context centering on the high-speed nature of chat interactions allowing for generally unacceptable exchanges (interpretations such as "this interaction is acceptable because it might have been that user 2 is responding to something a different user sent; user 2 didn't see user 1's message"), the directions also included a specification that 'All keysmashing users in the example contexts are responding directly to the provided messages or image; they are not responding to some other conversation in the same channel or chat.'

3.2.5. Contextless meanings of keysmashes

The sixth section contained only a grid question asking respondents to assign meanings to ten given keysmashes in a basic context ("Person A makes a statement or asks a question. Person B responds with the below keysmash. Person B is most strongly expressing...."). Only one option was permitted per keysmash to force respondents to choose the *strongest* meaning they interpreted from the item. Options for meanings included common emotions (anger, surprise, affection, amusement) and three discourse-related options (turn signaling, non-comprehension, attention signaling). The three remaning choices allowed respondents to indicate they did not know what the keysmasher meant, that they knew what they keysmasher meant but could not assign it to one of the given options, or that they felt the keysmash was unacceptable. These non-emotional response options were defined at the top of the grid question to ensure respondents understood what each option meant.

3.2.6. Survey transmission

The survey was intended to ask a wide audience a variety of questions in an easily transmissible way. An easily shared survey link was important both because the survey was essentially aimed at anyone who keysmashes online, and because the survey included questions about how different groups keysmash, so collecting data from several different groups was important. The survey was distributed via google forms; the survey link was shared on several discord.com chat groups and posted on Twitter, with encouragement to send the link to anyone who might be interested. The survey was opened for responses on November 26, 2019, and closed on May 11, 2020. The survey was shared and retweeted for a total of 1243 responses. No identifying data was collected, since knowledge of keysmashes was hypothesized to correlate with membership in potentially sensitive groups; the only demographic information collected was about the keysmasher's specific keysmash habits and those of their friends. As a result, there is little actual information about which respondents found the survey where; future surveys will include questions on where a respondent found the survey link, as well as potentially asking what social media platforms respondents spend most of their time on and where respondents mostly see keysmashes.

3.3. Experiment 1: Results & Discussion

As this experiment was conducted to investigate the linguistic nature of keysmashes, results are grouped according to relevant linguistic criteria: semanticity, arbitrariness, standards of form, convention, and traditional transmission.

3.3.1. Semanticity

Semanticity refers to a signal that intentionally indicates particular meanings (Hockett 1960). If keysmashes are properly semantic, they should convey meaning as a main purpose, rather than as a side effect of other occurrences. If keysmashes can convey meaning as a primary purpose, keysmashers should be able to understand what a keysmash means on its own, without context that might supply information about meaning.

Responses to the emotional grid in section six of the survey strongly indicate keysmash semanticity. Keysmashes were most frequently categorized as expressing amusement, surprise, or anger; no other options achieved majority on any question (and never exceeded 15% of responses for a particular item). The prevalence of particular emotional labels indicates that different, coherent, and categorizable emotions are expressed by different keysmashes. None of the ten keysmashes presented for the emotional grid were marked as unacceptable by a majority. Responses to some keysmashes were entirely mixed, where no response option gained more than 20% of responses; on the other hand, some keysmashes elicited particularly strong responses. 57% of respondents categorized the keysmash "sksksksksks" as indicating amusement; similarly, 55% of respondents categorized "LRGHGGGGH" as indicating anger. The keysmash "DBSJFHDKFJ" produced a notable tie, where 40% indicated surprise and 35% indicated amusement. While the strongest results seem only slightly above a majority, they are well above chance (as there were ten response options for each keysmash). The highest percentage response for each keysmash, even those with highly mixed responses, was above 20%.

Strong acceptability judgements on keysmashes in conversational contexts also indicate semanticity. Acceptability judgements in context should either indicate when it is acceptable to keysmash in a conversation, or conversely, exhibit a pattern where every keysmash is judged as not having a particular acceptability. Given that there are strong judgement results, keysmashes certainly mean *something*. While it is possible that keysmashes might be a form of expletive characterized by randomness (much akin to random shrieking in an emotional context), the emotional grid shows that keysmashes can express coherent emotions even outside of any given context. With this in mind, it is less likely that keysmashes are expressing random expletives (and thus less likely that the judgement results are largely about the acceptability of shrieking at various conversation points) and more likely that keysmashes are again representing particular meanings, with the judgement results indicating whether each keysmash's meaning is acceptable in that context. Keysmashes were rated on a scale of one to five, with one being least acceptable and five most acceptable. The averaged participant ratings of a particular keysmash will be referred to as its acceptability score.

Of particular note, similar keysmashes could be judged very differently depending on their context (or lack thereof): "HLSKJDFHLKJH" was judged to be very acceptable without context, with an acceptability score of 4.84. A keysmash composed of similar letters, "SDHKJSDHSD", was judged very unacceptable when used in response to a serious and emotionally charged situation, scoring 1.24. Keysmashes of various compositions that were presented in light conversational contexts, such as responses to jokes, were judged as very acceptable: "BJJKKDHKHKK", "djfdfdsjfggdtbnnhnhhf", and "hdsfgjbb" all had acceptability results over 3, with "hdsfgjbb" in particular rating 4.64 (compared to 3.31 and 4.43, respectively). These ratings are close to the acceptability ratings of other similar keysmashes presented for judgement without context: the keysmash "dhsnksdsdn" scored 4.40 and "asdiohuij" scored 3.79.

3.3.2. Arbitrariness & Standards of Form

Given the semanticity of keysmashes as discussed above, one may examine the forms of various keysmashes to draw conclusions about keysmashes' form-meaning relationships. Since keysmashes' forms appear to be random while still conveying coherent emotional meanings, keysmashes likely have an arbitrary form-meaning relationship.

However, if keysmashes had *completely* random forms, all keysmashes presented for contextless acceptability judgements should have received uniform results, since the form would be irrelevant to how people chose to rate keysmash acceptability without context to define any additional meaning. With standards of form, one would instead predict that characteristics of a keysmash would affect how participants rated the acceptability of keysmashes. Participant ratings of various keysmashes presented for contextless acceptability judgement did exhibit strong judgements on well-formedness. Keysmashes with characteristics such as excessive length, "special characters" that require pressing the shift key to type on standard QWERTY layouts, numbers, and several vowels were predicted to be judged unacceptable. Keysmashes in uniform case composed mostly of consonants found on the home row of a standard QWERTY keyboard were predicted to be judged acceptable.

Since judgements were made on a five-point scale, the standard deviation of responses was used to measure the variability of participant responses.

Two keysmashes received especially pronounced results: “HLSKJDFHLKJH” received an acceptability rating of 4.84 with a standard deviation of 0.49. “dhsnksdsdn” received a rating of 4.40 with a standard deviation of 0.95.The other eight keysmashes presented for contextless judgement averaged an acceptability rating of 2.70 with a standard deviation of 1.41. Five out of the total ten keysmashes presented for contextless judgement were chosen specifically to test the acceptability of keysmashes with particular properties, such as special characters, unusual lengths, or letters not located on the home row of a QWERTY keyboard. Low average scores and high variability among participants about these eight keysmashes were therefore in line with hypotheses predicting that these keysmashes would be judged unacceptable.

Keysmashes presented for acceptability judgement *with* context also provide evidence for standards of form. An unusual keysmash collected from chat logs and a similar, but invented, keysmash were presented in similar contexts for judgement. In the absence of standards of form, these keysmashes should have polled similarly. Instead, the resulting judgements differed both in response pattern and consensus: the unusual keysmash was clearly judged an unacceptable keysmash, with an acceptability rating of 1.94. The invented keysmash received very mixed results, with a nearly flat pattern of 20% per response across the scale. Its acceptability rating reflects this with a score of 3.06. There was no consensus about the invented keysmash’s acceptability, while the other keysmash was simply judged very unacceptable. The only difference between these keysmashes was a difference in form; participant sensitivity to this difference thus indicates standards of form.

3.3.3. Convention

Convention refers to group agreement on how meaningful forms might be used. Given that keysmashes are likely meaningful, if there are conventions about keysmash use, there should be differences in acceptability judgements for different uses of a keysmash.

Contextually based acceptability judgements and variability measures of those judgements indicate that there is general agreement not only about what keysmashes mean, but also about how they should be used. If there were no social agreement about keysmashes, participants should have judged the presented items either on the basis of their form or randomly; instead, the context in which items were presented changed how they were judged. Items for contextual judgements included the very similar keysmashes “dhdjaldhbfnznfnsfkjsfk” and “djfdfdsjfggdtbnnhnhhf”, presented in different contexts. The former keysmash was presented in a more unusual (but not uncommon) context, as an indication of the keysmasher’s mood. Participants were given the message below and asked to rate the acceptability of the keysmash in context.

> [User 1A] I’m running late on leaving the house and getting all “dhdjaldhbfnznfnsfkjsfk"

The latter keysmash was presented in a very common context, as a response to a joke posted by another user:

> [User 1A] us on halloween but also on every other day because this is just what our brain looks like actually

[User 2B] djfdfdsjfggdtbnnhnhhf

These keysmashes' forms are similar in length, starting letters, and general letter composition. Without social conventions about what keysmashes mean and how they can be used, these keysmashes should have scored very similarly, with the given contexts more or less irrelevant to the judgements. Instead, these keysmashes' acceptability scores and score variability among participants differed. The keysmash presented in the more unusual context scored 3.49 acceptability with a standard deviation of 1.29, while the keysmash presented in the more common context scored 4.43 acceptability with a standard deviation of 0.85. The latter keysmash was considered more acceptabl, with lower participant variability about the judgement, than the former. The acceptability ratings of these two keysmashes were significantly different ($t(2150)=-21.356$ p-value < 0.0005). Since these keysmashes had similar forms, this difference is driven by their uses in different contexts, indicating social convention about their use.

3.3.4. Traditional transmission

Traditional transmission involves learning the usage of language from others' use (as opposed to seemingly-innate knowledge of a language or language-like system as seen in animals). Anyone able to use a keyboard can produce random strings that might qualify as keysmashes, but this does not come with intuitions about acceptable forms and meanings. Intuitions about the use of keysmashes come from interactions with other keysmashers.

Participants who answered yes to the question "Do a majority of your friends keysmash often or very often?" had more granularity to their acceptability judgements, indicating community influence on keysmash-based intuition. This implies that participants who spend time around keysmashers develop finer intuitions about keysmashes' use and acceptability, leading them to make more granular decisions and use a greater range of the scale. Participation in a keysmashing community was more indicative of judgement granularity than self-identification as a keysmasher; participants who were not themselves keysmashers but whose friends keysmashed frequently made more granular judgements than those who self-identified as keysmashers but did not have friends who keysmash frequently.

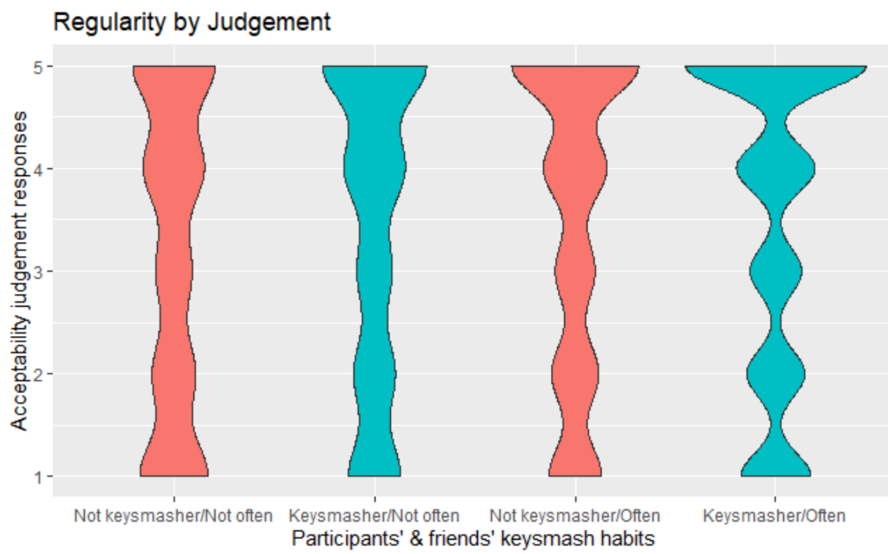

Figure 1. Granularity of keysmash judgements. Distributions in red are participants who did not identify as keysmashers, while distributions in blue are participants who identified themselves as keysmashers. The left two distributions answered that their friends did not keysmash regularly or often, while the right two distributions answered that their friends did.

Among their various uses, keysmashes are used to imply membership in or association with various groups. The Wikipedia entry for keysmashes notes that keysmashes are popularly associated with people in fandoms and LGBTQ+ people. This appears not to be coincidental; several participants responded directly with variations of "because i'm gay" in the "other" option of the question asking why one can keysmash.

Participants were asked to rate how likely members of various groups are to keysmash on a scale of 1 to 5, where 1 was less likely and 5 was more likely to keysmash. Groups were rated independently of each other. Out of nine groups, people in fandoms were rated most likely to keysmash, with an average of 4.653 and standard deviation of 0.649. LGBTQ+ people were second most likely to keysmash, averaging 4.485 with a standard deviation of 0.775. People under the age of 25 had very similar results, with an average of 4.423 and standard deviation of 0.788. Participant ratings of these groups support the popular association of keysmashes with younger people, LGBTQ+ people, and people in fandoms.

Participants were then asked whether they feel they can recover social information about keysmashers based on their keysmashes. If keysmashes are used to signal information about the keysmashers, others should be able to pick up on this information. Participants' ability to recover social information was above chance, with 61.8% of respondents answering that they could. 20.6% of participants indicated that their ideas were almost always accurate, while 29.0% indicated that they were no more likely than chance. 12.1% responded that their intuitions about a person based on their keysmashes were often stereotypes that the keysmasher in question did not always fit.

Participants' perceived accuracy is slightly less relevant than the fact that the majority of participants do recover social information from keysmashes. In particular, recovering stereotypical social information from keysmashes indicates that there *is* a stereotypical keysmasher, whether a given keysmasher fits it or not. That more participants self-reported accuracy than not could indicate that keysmashes do convey a decent amount of information about keysmashers, or that enough keysmashers fit the stereotype that intuitions based on said stereotype are accurate.

3.5. Experiment 1: Conclusion

While certainly not a language unto themselves, based on the results of this survey, keysmashes do appear to behave linguistically. As described above, keysmashes convey meaning, have standards of form, have arbitrary relationships between those forms and meanings, and involve social conventions for their use. Keysmashers have intuitions about well-formed and ill-formed keysmashes, find keysmashes meaningful both with and outside of context, and use keysmashes to imply their association with particular social groups. Keysmash use is traditionally transmitted: not anyone who can type a keysmash necessarily knows how it is used or what it means in communities where people frequently keysmash. Investigating keysmashes as linguistic phenomena indicates that at least some Internet-based unorthodox communication works in ways analyzable with, and relevant to, traditional linguistic methods and tools. By these familiar means, one may conclude that keysmashes are not entirely random, nor are they meaningless.

4. Experiment 2

4.1. Experiment 2: Introduction

While keysmashes generally appear random, Experiment 1 instead found keysmashes to have various properties characteristic of language; perhaps most notably, keysmashes exhibit semanticity, arbitariness, and standards of form. Experiment 2 focused on the standards of form, and circulated a survey to examine the well-formedness of keysmashes with a constraint-based analysis. A total of eight primary constraints were hypothesized to have an effect on keysmash well-formedness: length, keyset, punctuation, vowel, lexical, phonotactic, repetition, and gesture. Constraints such as Punctuation or Vowel are violated when instances of what they are named for exist in a keysmash. Constraints such as Length were hypothesized to function in a sort of bell-curve: having too little or too much of the described quality would incur violations.

To investigate these hypothesized constraints, we created a survey that primarily asked participants to make acceptability judgements. Results were analyzed with particular attention to the existence and size of effect of the hypothesized constraints, utilizing linear models and aggregates of average acceptability ratings; discussion of analysis follows.

4.2. Experiment 2: Methods

Hypothesized constraints on keysmash well-formedness were examined with a survey consisting primarily of acceptability judgements on keysmashes designed for this purpose.

4.2.1. Survey Design

The survey was divided into three sections. The first section of the survey collected basic information about the participant. Participants were asked whether they knew what a keysmash was prior to taking the survey, how often they keysmash, and how often their friends keysmash. (The last question was included as the previous experiment showed a community effect on the judgement of keysmashes, affected by both whether the participant themself keysmashes and whether their friends keysmash.) The current study was interested in both replicating this finding and investigating whether various communities rank constraints differently; that is, whether different communities have different senses of keysmash well-formedness. In the same vein, participants were asked what social media platform they use most, then asked to rank their top three social media platforms by usage, out of six provided platform options. (Choices for "Other", "I do not use this platform", and "I use this platform, but not as part of my top three", were also included.) Finally, participants were asked whether keysmashes give them ideas about the person who typed them: specifically, whether seeing keysmashes gives an impression of what group the writer is from, and what kind of social media platforms they use. Keysmashers' potential ability to gather this information from reading keysmashes would suggest that there are community-specific standards of keysmash well-formedness.

The second section of the survey elicited keysmashes from participants who indicated they keysmash. The form was created such that participants who responded that they only keysmash once in a while, or never keysmash, were not shown this section. These participants were automatically moved ahead to section three, acceptability judgements, after completing section one. Section two provided a variety of contexts with instructions to keysmash in the provided short answer box if one would keysmash in response to the context. Participants were instructed to paste in "none" if the participant ordinarily keysmashes, but would not keysmash in response to that context. In case a participant who did not ordinarily keysmash was accidentally shown this section, participants were also instructed to paste in n/a for every context if one does not keysmash at all. As is the nature of participants given several blank spaces to respond to, the elicitation did not actually create neatly coded data, and many participants typed out actual responses to the contexts provided rather than paste "none". Although these data were collected, we will not analyze them here, as the major focus of the project was on the characteristics of keysmashes, which were probed by section three.

Section three presented 47 keysmashes for acceptability judgement (48 were designed, but one was accidentally excluded during survey creation). In contrast to Experiment 1, where the majority of keysmashes presented for judgement were collected from actual online "utterances" from public chatrooms, the keysmashes in this survey were designed to test the existence and influence of various constraints on keysmash well-formedness.

4.2.2. Creating the survey stimuli

Keysmashes were designed by taking "base" keysmashes, noting their constraint violations, then transforming them to repair the violation, introduce other violations, or otherwise change the relevant constraints. The transformed keysmash was then transformed again by the same method to create sets of three related keysmashes. Sixteen keysmashes were selected as "base" keysmashes and transformed twice to create the total of 48 keysmashes designed for acceptability judgements. Where possible, the base keysmashes were collected from public chat groups; however, in some cases, the base keysmashes were created by the author, to facilitate comparison and transformations of keysmashes with constraint violations not often attested.

The constraints were defined, and their related transformations performed, as follows:

4.2.3. The 'length' constraint

Keysmashes that are either too long or too short violate the length constraint. Keysmashes that fall within the range of 5 to 20 characters were hypothesized to be more acceptable than those outside it, based on brief reference to keysmashers' intuitions. Keysmashes incur one violation for every five characters off from the minimum or maximum amounts. For example, a keysmash with 21 characters and a keysmash with 25 characters would both have only one length violation. Any keysmashes with only one, two, three, or four characters would only have one violation. (While keysmashes with such few characters should be very unacceptable, none were included in the stimuli, as the current investigation was more focused on effects of longer keysmashes, not shorter; in more detailed examinations where very short keysmashes are included in the stimuli, the definition of length should perhaps be revised to properly penalize shorter keysmashes.)

To correct length violations, the *center* X number of characters (where X was approximately half the length of the keysmash) were removed. As characters from the beginning and the end of a keysmash were hypothesized to be meaningful, those characters could not simply be removed without possibly changing the meaning of the keysmash. (The question of what *meaning* change might result was not an integral part of keysmash *well-formedness*, and was not addressed directly in this study; however, having keysmashes in transformed sets as similar to each other as possible aside from constraint-based changes was preferable.)

4.2.4. The 'keyset' constraint

The character makeup of a keysmash was hypothesized to have an effect on its acceptability. This constraint is referred to as "Keyset". In Experiment 1, keysmashes that include punctuation were rated less acceptable than those only containing letters from the "home row" of a QWERTY American English keyboard, which were generally rated very acceptable. As a constraint, "Keyset" is violated when keysmashes include characters that are not on this home row. Punctuation, including the semicolon and quotation mark; numbers; and "special symbols" such as @, #, or % all constitute violations of Keyset, along with all letters that are not on the home row.

A touch-typing finger placement chart was particularly useful for creating and correcting violations of keyset. The chart shows the ideal placement of each hand, color-coding which finger should be used to strike any given key. To create violations of keyset, home row characters in the base keysmash were replaced with characters typed with the same finger on a different row on the keyboard. To correct violations of keyset, characters off the home row were replaced with characters on the home row in approximately the same way.

Figure 2. A diagram indicating what fingers are to be used to strike each key when touch-typing on a typical QWERTY American English keyboard. (Bolte, 2021.)

4.2.5. The 'punctuation' constraint

Keysmashes that include punctuation violate the 'punctuation' constraint. Keysmashes that included punctuation in Experiment 1 were rated less acceptable than keysmashes without punctuation. As this constraint's existence had support from Experiment 1, and other constraints hypothesized had several other directions to explore, this constraint was not the focus of this study, and only one set of keysmashes was transformed concerning this constraint. (Excluding this constraint entirely would have painted a less complete picture of keysmash well-formedness.) To create violations of punctuation, some characters were replaced with punctuation, again maintaining the finger used for the keystroke according to the typing chart. To correct violations of punctuation, instances of punctuation were *deleted*. This does not parallel the method used to create violations, but in the single instance where a violation of punctuation was corrected as part of a transformation, the only possible replacement would have been "p", which incurs its own violation for keyset. The offending semicolon was therefore deleted.

Keysmashes that include punctuation will violate *both* punctuation and keyset. While Keyset could have been defined such that it is only violated by letters that are not on the home row, and separate constraints proposed (but excluded) for numbers and special characters, this seemed unnecessarily complex for the present study. The stimuli did not include any keysmashes transformed to include special characters or numbers, as Experiment 1 had established that keysmashes which include these are judged strongly unacceptable; including keysmashes with such characteristics would have therefore created a more complex set of constraints perhaps only for complexity's sake. The current investigation was, however, interested in the effect of punctuation over the other types of keyset violations, and therefore included Punctuation as its own specific constraint.

4.2.6. The 'vowel' constraint

Having a vowel in the keysmash was hypothesized to render it less acceptable. This constraint is referred to as "Vowel". To correct violations of this constraint, the vowels in the keysmash were either replaced with nearby homerow keys according to the fingerings on a typing chart created by the typing academy (Bolte 2021) or deleted outright. The two different methods of vowel correction allowed us to also examine differences in the effects of the gesture and vowel constraints (which are elaborated on in analysis). As a good number of keysmashes collected from public chat contexts featured vowels, this constraint was not intentionally created during any transformations; vowel constraint violations were accounted for in "base" keysmashes.

4.2.7. The 'lexical' and 'phonotactic' constraints

The constraint 'lexical' is violated when a keysmash contains a word recognizable as english. ('Lexical' is considered and coded separately from 'phonotactic', which is violated when a keysmash contains a string that follows the rules of English phonotactics, but is not a word in and of itself. Every violation of lexical is also a violation of phonotactic.) English was the only language considered for constraints on keysmash form because English was the assumed shared language of participants taking the survey; other language data would have to be collected from other surveys that specified participants who would know those languages. Violations of lexical were created by inserting vowels into consonant shapes that already existed in base keysmashes, or by changing single existing characters to create words. To correct violations of lexical, vowels were moved or deleted according to the corrections for vowel violations described above.

4.2.8. The 'repetition' constraint.

Repetition was a difficult constraint to define. It was largely included out of keysmashers' intuitions that a particular amount of repetition positively affects the acceptability rating of a keysmash, while either having no repetition or being composed entirely of identically repeating characters negatively affects the acceptability of a keysmash. For purposes of categorizing keysmashes during design and transformation, repetition was hypothesized to exist where obvious repetitions of substrings existed in base keysmashes (the repetition of two or three, or more characters, next to each other). For purposes of analysis, this simple definition was replaced with n-gram and skip-gram counts, described in detail in section 4.6, analysis of the repetition constraint.

As this constraint was difficult to define, violations of repetition were largely not created or corrected in the keysmash transformations. In a single case, a violation of repetition was corrected by inserting characters whose positions are not adjacent to the repeating characters' on the standard qwerty keyboard. (These characters were all home-row letters, and placement alternated left-hand and right-hand characters, so as not to violate other constraints while correcting the violation of repetition.)

4.2.9. The 'gesture' constraint

Keysmashes that are formed by performing a certain keyboard gesture were hypothesized to be more acceptable than those formed with other ones. This acceptable gesture involves the use of both hands on or near the home row, alternating hands, performing two to three keystrokes with each hand before switching over. Which keys are on which side of the keyboard were again assigned according to the touch typing chart. This constraint was formulated in response to keysmasher intuitions that some keysmashes which did not violate any of the other constraints were more acceptable than others, implying the existence of an additional constraint beyond the other seven originally hypothesized.

There were two methods used to create violations of this constraint. In the first method, characters in the base keysmash were moved so that the resulting keysmash was composed of characters all on one side of the keyboard, then the other. In the second method, which was only used once, a gesture violation was created out from a base keysmash that was composed only of two repeating characters. The violation was created by inserting home row characters that alternated keyboard sides.

To correct violations of gesture, characters from opposite sides of the keyboard were paired together, then arranged among the other groups to create a keysmash that a typical touch typist would use both hands, alternating, to type.

Set Number	Base keysmash	Hypothesis	Transformation	Hypothesis I	Transformation II	Hypothesis II
1	erdtfyguyh	vowel	drdtfygjyh	keyset	dfdgfhgjhh	acceptable
2	asdfsdfasjfas	vowel, repetition	sdfsdfsjfs	repetition	sdfsjsdfs	acceptable
3	hdhsdsfj	acceptable	ydysesrj	vowel, keyset	ydysdsrj	keyset

4	dkhasfjh	acceptable	dkhasofjh	lexical, vowel	dinkhasfontjh	lexical
5	shshdhdh	repetition, acceptable	shdsjhfdkhsdjh	gesture	sdhfsdhjkljh	acceptable
6	snsnabehabsja	lexical, vowel	snsnbdhbsj	gesture	snsndbshjbs	repetition, acceptable
7	kjdngilsejrhu	keyset, vowel	kjdngklsdjrhj	gesture	kdjgnksldjhrj	keyset, acceptable
8	oirqewuq	vowel, keyset	lkrqdwjq	keyset	lkfadsja	lexical, acceptable
9	ssdiulfsodylfkjsdghsjkafhdjsklgsalfdshnfuslhjkl	length, vowel, keyset	ssdlfsdylfkjsdghsjkfhdjsklgslfdshnfslhjkl	length, keyset	ssdlfsdylfkshnfslhjkl	repetition
10	SDLKJAFLS	acceptable	SCLKMZVLS	keyset	SCL,JZF.S	punctuation, keyset
11	LKSDJFLKAJSFD	acceptable	LKXDMFKZJXFD	keyset	LKSADJFILKAJESFD	lexical, vowel
12	SJDKFSKDLNGKSLDNGSDKLGSKDLGNDSK	length	SJDKFSKDLNGSKDLGNDSK	length, acceptable	NGKSLDNDKLGSKDLGNDSK	keyset, acceptable
13	KSNCKWNDOWNFONELDNFOENR	length, lexical	KSNCKWNDLWNFLNDLDNFLDNR	keyset, length, repetition	KSJFKSJDLSJFLJDLDJFLDJF	repetition, acceptable
14	fhaskjfd	lexical, vowel	fhskjfd	acceptable	fdsfhkj	gesture
15	LKFJSAOILKSJDF	lexical, vowel	LKFJSAMILKSJDF	lexical	LKFJSAJKLKSJDF	repetition, acceptable
16	EALW;IFDJKSLA	vowel, punctuation	LWFDJKSL	keyset	LSWFDJKL	gesture, keyset

Table 1. Keysmashes according to transformation set. Primary hypotheses for each keysmash are listed to the right of the keysmash. Note that not all violations of constraints are listed beside each keysmash, in this table; constraints were listed according to major hypothesis.

	Length	**Punctuation**	**Vowel**	**nHomerowVowel**	**Keyset**	**Lexical**	**Phonotactic**
sdfsdfsjfs	0	0	0	0	0	0	0
LKFJSAOILKSJDF	0	0	3	2	2	0	1
fdsfhkj	0	0	0	0	0	0	0
SJDKFSKDLNGSKDLGNDSK	0	0	0	0	0	0	0
dkhasofjh	0	0	2	1	0	1	1
fhaskjfd	0	0	1	0	0	1	1
SCLKMZVLS	0	0	0	0	4	0	0

sdhgjksdgndf	0	0	0	0	0	0	0
KSNCKWNDLWNFLNDLDNFLDNR	1	0	0	0	10	0	0
sdfsjsdfs	0	0	0	0	0	0	0
EALW;IFDJKSLA	0	1	3	1	3	0	0
ydysesrj	0	0	1	1	3	0	0
LKFJSAJKLKSJDF	0	0	1	0	0	0	0
dinkhasfontjh	0	0	3	2	5	3	3
erdtfyguyh	0	0	2	2	5	1	1
LKFJSAMILKSJDF	0	0	2	1	2	2	2
hdhsdsfj	0	0	0	0	0	0	0
kdjgnksldjhrj	0	0	0	0	0	0	0
LKSDJFLKAJSFD	0	0	1	0	0	0	0
asdfsdfasjfas	0	0	3	0	0	0	2
KSNCKWNDOWNFONELDNFOENR	1	0	5	5	8	1	3
lkrqdwjq	0	0	0	0	4	0	0
LWFDJKSL	0	0	0	0	1	0	0
NGKSLDNDKLGSKDLGNDSK	0	0	0	0	4	0	0
snsndbshjbs	0	0	0	0	4	0	0
LKSADJFILKAJESFD	0	0	4	2	2	0	1
ssdiulfsodylfkjsdghsjkafhdjsklgsalfdshnfuslhjkl	5	0	6	4	5	1	2
ydysdsrj	0	0	0	0	3	0	0
LSWFDJKL	0	0	0	0	1	0	0
snsnbdhbsj	0	0	0	0	4	0	0
drdtfygjyh	0	0	0	0	4	0	0
shdsjhfdkhsdjh	0	0	0	0	0	0	0
fhskjfd	0	0	0	0	0	0	0
SDLKJAFLS	0	0	1	0	0	0	0
dfdgfhgjhh	0	0	0	0	0	0	0
LKXDMFKZJXFD	0	0	0	0	4	0	0
ssdlfsdylfkjsdghsjkfhdjsklgslfdshnfslhjkl	4	0	0	0	2	0	0
oirqewuq	0	0	4	4	8	0	0
SJDKFSKDLNGKSLDNGSDKLGSKDLGNDSK	3	0	0	0	3	0	0
kjdngklsdjrhj	0	0	0	0	2	0	0
snsnabehabsja	0	0	4	1	5	1	2
dkhasfjh	0	0	1	0	0	1	1

SCL,JZF.S	0	2	0	0	4	0	0
kjdngilsejrhu	0	0	3	3	5	0	0
shshdhdh	0	0	0	0	0	0	0
lkfadsja	0	0	2	0	0	1	1
ssdlfsdylfkshnfslhjkl	1	0	0	0	2	0	0

Table 2. Keysmashes as they were coded for six of the eight major violations (not including gesture or repetition, but including the constraint non-homerow vowel, reasons for which are described in analysis, section 4.6.)

4.2.10. Survey Administration

The survey, approved under UCSD IRB #210159, was available using Google Forms from February 2021 until April 2021, collecting 1039 responses over the period of two months. It was posted in various public chat groups hosted on Discord and tweeted on Twitter, with a request to spread the survey to friends after it was taken. A survey question asking where the respondent found the survey link revealed that the majority of respondents did find the survey on these two platforms; only 7.9% of participants found the survey through other means, amounting to 82 people out of 1039 who did not come from Twitter or Discord.

4.3. Experiment 2: Analysis

Analysis focused on the eight major hypothesized constraints: vowel, lexical, phonotactic, keyset, punctuation, length, repetition, and gesture.

Once the survey was closed, results were tabulated using R. Constraints were examined using linear regression models in R in 'lm' to determine statistical significance. Every major hypothesized constraints (those listed above) resulted in a significant $p < 0.005$. (There were also other constraints hypothesized and examined, some of which had different p-values when analyzed; these constraints are included in the table below, and are discussed in their relevant sections). The potentially surprisingly high number of significant results is understandable given that there were more than a thousand participants.

Constraint name	Coefficient	p-value
keyset	-0.212	< 0.005 ***
punctuation	-0.523	< 0.005 ***
length	-0.148	< 0.005 ***

lexical	-0.284	< 0.005 ***
phonotactic	-0.200	< 0.005 ***
vowel	-0.143	< 0.005 ***
non-homerow vowel	-0.253	< 0.005 ***
gesture	0.173	< 0.005 ***
repetition-boolean	0.133	< 0.005 ***
n-gram count	0.008	< 0.005 ***
skip-gram count	-0.0006	< 0.005 ***
large skip-gram count	-0.0008	= 0.200 (n.s)
adjusted total of large skip-gram count	-0.0007	= 0.194 (n.s)

Table 3. Coefficients and significance of each constraint predicting acceptability.

The major hypothesized constraints had p-values indicating significance. Notably, some elaborated attempts to code repetition with skip-grams were not significant. (The meaning of, and reasoning behind, these constraints and their coding is discussed in the section on repetition, 4.6.)

With significance of effect generally established, this investigation turned its focus to the size of effect of each constraint.

4.3.1. Analysis of the Constraints Vowel and Non-homerow Vowel

Vowel is violated when keysmashes contain vowels. This constraint, however, penalizes the vowel "a", which is located on the home row. The home row was generally hypothesized to be the most acceptable set of characters for a keysmash (as captured in the constraint "Keyset"). The constraint "non-homerow vowel" was therefore created to investigate whether a vowel's placement on the keyboard matters with respect to its potential effect on a keysmash's acceptability. 'Non-homerow vowel' draws a distinction between vowels that do not incur keyset violations and those that do by not penalizing "a".

Due to the placement of most vowels on the keyboard, violations of vowel are often reflected in keyset, penalizing keysmashes with any vowels other than "a" twice. As this study did not perform ranking or weighting analysis of the constraints, and instead focused on size of effect, this double penalty was not particularly concerning. Interest in the potential difference of size of effect between vowel and non-homerow vowel themselves, without an interaction of keyset, motivated the inclusion of non-homerow vowel as a unique constraint.

Linear models were run predicting the acceptability of keysmashes with regards to the constraints vowel and non-homerow vowel. Both models showed a significant ($p < 0.005$) effect. The models indicate that there is a difference between vowels on the home row ("a") and those that are not; the presence of non-homerow vowels (coefficient of -0.253 per violation) was sworse than the presence of any vowel (coefficient of -0.143 per violation). There were also differences in average rating per violation:

Vowel violations	**Average rating**
0	3.068
1	3.167
2	2.872
3	2.801
4	2.423
5	1.929
6	2.234

Non-homerow vowel violations	**Average rating**
0	3.135
1	2.600
2	2.672
3	2.508
4	2.080
5	1.929

Table 4. Averages of vowel and non-homerow vowel constraints. (Note that non-homerow vowel only ever incurs up to five violations in a single keysmash, while vowel goes up to six.)

The presence of vowels that are not from the home row generally render a keysmash less acceptable than the presence of any type of vowel. Interestingly, the average rating of vowel *increases* from 5 to 6 vowels (Table 4). This was hypothesized to be an effect of length, where a longer keysmash is permitted more vowels.

To evaluate whether there is a meaningful effect of length on vowel acceptability, a model comparison was performed between a model including both length and vowel versus a model where vowel and length interact. While both constraints and the interaction were significant ($p < 0.005$) in the interacted model, the ANOVA indicated that the interaction actually decreases the overall quality of fit over the model that simply includes both constraints ($p < 0.005$). In the interacted model, vowel and length had negative coefficients, as expected (-0.160 and -0.182 respectively), and the interaction did move the coefficient in a positive direction (0.032), but not in such a way that the interaction created a better-fitting model. This indicates that the slight increase in acceptability when comparing keysmashes with five and six violations is not a result of interactions with length. The slight increase may be better explained by noting that the stimuli used in this study were transformed with an eye towards creating comparable sets, rather than comparable violations of individual constraints; keysmashes with five vowel violations may have coincidentally included worse violations of other constraints than keysmashes with six vowel violations. This explanation would also account for the relatively small rise, and the fact that it only occurs in the change from five to six vowel violations, rather than occurring over a larger portion of higher vowel violations.

As predicted, vowel had a generally negative effect on keysmash acceptability, and non-homerow vowel had a stronger negative effect on keysmash acceptability. This implies the existence of a constraint against vowels on keysmash well-formedness that functions as hypothesized.

4.3.2. Analysis of the Constraints Lexical & Phonotactic

Lexical is violated when keysmashes contain English words. Similarly, phonotactic is violated when keysmashes contain strings that follow English phonotactics.

Linear models were run predicting the acceptability of keysmashes with regards to the constraints lexical and phonotactic. Both models showed a significant ($p < 0.005$) effect. The models indicate that there is a slight difference between lexical and phonotactic effects: lexical violations (coefficient of -0.284 per violation) were worse than phonotactic violations (coefficient of -0.201 per violation). There were also differences in average rating per violation:

Lexical violations	Average rating
0	3.031
1	2.727
2	3.111
3	1.811

Phonotactic violations	Average rating
0	3.000
1	3.114
2	2.834
3	1.870

Table 5. Aggregates of vowel and non-homerow vowel constraints, by average acceptability rating.

From the tables, there is not a directly obvious relationship between these constraints. This is somewhat surprising, as every violation of lexical is by definition also a violation of phonotactic; it was hypothesized that one would be clearly worse than the other, but this is not immediately obvious from the compared ratings.

Given that the relationship between these lexical and phonotactic constraints and acceptability was non-linear; that there are clear links to other, more clearly defined constraints such as vowel (and, presumably, non-homerow vowel); and that the dataset did not include a plurality of items to evaluate these hypotheses; further evaluation of these constraints will be reserved for future work.

4.3.4. Analysis of the Punctuation Constraint

Punctuation is violated when keysmashes contain punctuation. In Experiment 1, keysmashes containing punctuation, numbers, and special characters (those that necessitate holding down the shift key to type, such as #, $, or %) were universally judged less acceptable than those without. In the present study, the punctuation constraint was included to create a fuller picture, but without as much focus as other major hypotheses, as results of Experiment 1 implied a very strong negative effect when *any* non-letter characters are present in a keysmash. As a result, no additional constraints were hypothesized to differentiate between punctuation considered part of the home row (";") and other punctuation, or to differentiate between punctuation, numbers, and special characters. This presents an opportunity for future work and improvement.

A linear model was run predicting the acceptability of keysmashes with regards to the punctuation constraint. The model showed a significant ($p < 0.005$) effect. The coefficient for punctuation, at -0.523, had the greatest absolute value among the constraints tested in this study, indicating that punctuation has a relatively strong effect on keysmash acceptability. These results are somewhat more impressive for their statistical significance, as there were only two keysmashes with punctuation violations included in the stimuli. The first keysmash, "SCL,JZF.S", was judged very clearly unacceptable, with an average rating of 1.82 (on a scale of 1 to 5, where 1 was least and 5 was most acceptable). "EALW;IFDJKSLA: was judged slightly more acceptable, with an average rating of 2.70.

Despite the sample size, punctuation had a clearly negative effect on keysmash acceptability.

4.3.5. Analysis of the Keyset constraint

Keyset is violated when keysmashes contain anything other than a letter from the home row. (Every violation of punctuation and non-homerow vowel is also a keyset violation.)

A linear model was run predicting the acceptability of keysmashes with regards to the keyset constraint. The model showed a significant ($p < 0.005$) effect. The coefficient for keyset was -0.212.

Keyset violations	Average rating
0	3.571
1	2.822
2	3.003
3	2.463
4	2.537
5	2.257
8	1.883
10	2.152

Table 6. Aggregate of keyset violations with average acceptability rating. Note that no individual keysmash had 6, 7, or 9 violations of keyset.

As would be indicated by the negative coefficient, acceptability decreases as violations increase. However, this is not a linear effect, especially considering the rise in acceptability when moving from 8 to 10 keyset violations. This rise could be explained by a relationship with length where longer keysmashes are allowed more violations. To test this, two models were constructed for comparison. The first model included both length and keyset, while the second included an interaction between the two constraints. The models were compared using ANOVA, which determined that the interaction does improve overall fit. In the interacted model, both constraints and the interaction were significant ($p < 0.005$). Coefficients for length (-0.303) and keyset (-0.233) were in the expected negative direction; the keyset*length interaction was in a positive direction (0.075). This lends support to the idea that longer keysmashes made keyset violations more palatable to participants, and stands in contrast to the finding that longer keysmashes do not make *vowel* violations more palatable to participants.

As predicted, keyset had a generally negative effect on keysmash acceptability, with the additional interesting finding that longer keysmashes are still considered acceptable with some additional amount of keyset violation.

4.3.6. Analysis of the Length Constraint

Length is violated when keysmashes are either too short or too long. Keysmashes were coded for length violations according to their character count: keysmashes with a length of five to twenty characters incurred no violations, as referenced from keysmasher intuitions about acceptable lengths. Every five additional characters were counted as one violation, such that keysmashes with 21 characters and 25 characters would both be penalized once, while a keysmash with 26 characters would be penalized twice. In this study, the longest keysmash in the stimuli had a length of 47, and incurred five length violations. (Exceptionally short keysmashes were not included or coded for; no keysmashes among the stimuli were shorter than seven characters.)

A linear model was run predicting the acceptability of keysmashes with regards to the keyset constraint. The model showed a significant ($p < 0.005$) effect. The coefficient for keyset was -0.148, which is relatively small compared to the coefficients of other constraints. This may be because the average acceptability is highest when a keysmash has three length violations, rather than patterning linear behavior.

Length violations	Average rating
0	3.028
1	2.346
3	2.919
4	2.445
5	2.324

Table 7. Aggregate of length constraint with average acceptability rating. Coincidentally, no keysmashes of length 26 to 30 were tested, so there are no results for keysmashes with two length violations.

A keysmash with length 5 - 20 (coded as '0' above) or 31 - 35 (coded as 3 above DO THE REST) was more acceptable than keysmashes with character lengths of 21 to 25. It was unintentional that there were no keysmashes with lengths 26 to 30. The presence of stimuli with this length may have helped identify exactly how sensitive keysmashers are to keysmash length: the relatively pronounced differences in average rating between the various violations of length are very interesting, as only five (or fewer) characters separate one length violation from the next.

Length violations had interestingly nonlinear effects on keysmash acceptability.

4.3.7. Analysis of the Repetition Constraint

Repetition, and what constituted a violation of repetition, was difficult to define. Initially, repetition was coded with one violation when a substring of two or more characters repeated in a keysmash. While this definition is straightforward, it did not always match keysmashers' intuitions about when instances of repetition actually happen. When repetition is coded this way, the keysmash “hdhsdsfj” would not score any violations of repetition (as no substrings within the keysmash repeat). However, keysmashers reported thinking of this keysmash as repeating two or three “pieces” or “chunks”, which translated for the purposes of this study as incurring two or three violations of repetition. Since the definition of repetition as repeating substrings could not account for keysmasher intuitions on violations of repetition, but the intuition that something repeated stood, we turned to skip-grams. Skip-grams count the repetition of characters or words that are not necessarily directly adjacent; in this case, with two possible characters’ distance between items. The same keysmash above had 16 potential strings of length 2 to 5 that a simple Python script, using NLTK for n-gram and skip-gram analysis (with a skip-gram window of 2), could detect as having at least two repetitions.

Substring	Repetition count
ds	2
hd	2
hs	3
sf	2
dhs	2
hds	2
hsd	2
hsf	2
sfj	2
dhsf	2
hdhs	2
hsds	2
hsfj	2
dhsfj	2
hdhsf	2
hsdsf	2

Table 8. Skip-gram analysis of “hdhsdsfj”, showing possible repeating substrings.

As the skip-gram approach tended to note relatively high numbers (the keysmash "SJDKFSKDLNGKSLDNGSDKLGSKDLGNDSK" had the highest number of detected repetitions, at 288), the results of the skip-gram script were coded in two additional ways: once counting only "large" skip-grams, where only detected substrings with length four or five were counted, and once adjusting this total for large skip-grams, by adding violations each time a large skip-gram is detected in the keysmash more than twice. These modifications were analyzed alongside the major hypothesized constraints as "large skip-gram count" and "adjusted total of large skip-gram count". The initial definition of repetition, though it did not fit keysmasher intuitions about repetition, was included for the sake of comparison, under the name "n-gram count" (as n-gram is another name for the first, simple definition of repetition). If a keysmash had any skip-grams at all, it was coded as having an instance of "repetition boolean", which indicates whether a keysmash does or does not have skip-grams without regard for how *many* skip-grams it may have. Results regarding the p-values and coefficients of these various definitions of repetition can be seen in Table 3 beside the major hypothesized constraints.

Approaches that modified skip-gram counts did not yield statistical success, though the simpler approaches of "repetition boolean", "n-gram count", and "skip-gram count" were shown to be statistically significant.

Linear models were run for each method of coding repetition, predicting the acceptability of keysmashes with regards to these constraint definitions.

Model Name	**Coefficient**	**p-value**
repetition-boolean	0.133	< 0.005 ***
n-gram count	0.008	< 0.005 ***
skip-gram count	-0.001	< 0.005 ***
large skip-gram count	-0.001	= 0.200 (n.s.)
adjusted large skip-gram total	-0.001	= 0.194 (n.s.)

Table 9. Coefficients and p-values of the linear models run for each analyzed definition of repetition.

The lack of significance found for both the models of large skip-gram count and the adjusted total of large skip-gram count suggest that the repetition constraint is not reliant on the size of the repeating substring. While the model for n-gram count was found to have a significant effect, the definition of n-gram count does not fit with the author's keysmash-related intuitions on the definition of repetition. The model of skip-gram count has a significant effect, and fits keysmasher intuitions, but has a very *small* effect. This leaves the model of repetition-boolean as the relative best definition of repetition among those tested, with a significant effect that does not clash with keysmashers' intuitions, and with a somewhat larger coefficient, indicating larger effect. However, this constraint functions differently than the hypothesized effect of repetition: its coefficient is positive, while repetition was hypothesized to have a negative effect.

Repetition was difficult to define, and as such a variety of possible definitions were analyzed; after creating linear models according to each possible definition of repetition, the relative best definition was repetition-boolean, which only measures whether a keysmash has skigrams or not.

4.3.8. Analysis of the Gesture Constraint

Gesture is violated when keysmash letter composition is indicative of a gesture that did not involve alternating both hands on the keyboard. Gesture was coded in a series of steps. First, keys were assigned "left" or "right", according to the typing chart (Bolte 2021). Each keysmash was then translated into a form that represented only the use of the left or right hand to type, with 0 representing right and 1 representing left. This result was averaged by the sum of its digits to create a basic quantitative measure of whether the original keysmash was comprised of an *even amount* of left or right keystrokes. (This method of coding gesture ignores completely whether the hands actually alternate during the typing process, which was hypothesized as important to the gesture constraint. However, devising a method to code this in a quantifiably analyzable way proved to be quite difficult.) Based on this "gesture average", it was determined whether the original keysmash had any gesture violations or not: keysmashes with gesture averages outside of the "acceptable" range of 0.4 to 0.6 were coded as having gesture violations. (This range was centered on 0.5, which is the gesture average a perfectly left-right balanced keysmash would have. The range includes 0.1 deviation in either direction.) Keysmashes scored an additional violation for every additional 0.1 deviation away from the acceptable range.

Keysmash	Gesture (left-right)	Gesture average	Gesture violations
kdjgnksldjhrj	0101001010010	0.385	1
fhaskjfd	10110011	0.625	0
sdfsjsdfs	101101111	0.778	2
asdfsdfasjfas	1111111110111	0.923	4

Table 10. Examples of gesture coding.

Gesture violations	Average rating
0	2.911
1	2.902
2	3.004
3	4.178
4	3.620

Table 11. Aggregate average acceptability ratings by gesture constraint.

Similarly to repetition, gesture did not function as it was hypothesized to work: keysmashes with three or four gesture violations were rated *more* acceptable than those with fewer. This may be for a variety of reasons. Keysmashers may simply find keysmashes that are “unbalanced” with regard to left and right keys more acceptable, or the acceptable range for gesture coding may have been too small. Notably, this method of coding gesture did not account for the actual alternation of hands, only quantifying the total use of keys assigned to the left and right hands; the ability to account for this factor may affect future results.

The gesture constraint did not function as hypothesized.

4.4. Experiment 2: Discussion

Punctuation, length, vowel, and keyset had fairly clear and easily interpreted effects on keysmash acceptability. Lexical and Phonotactic had interpretable effects, though these effects could not be examined in the depth they may have indicated. Repetition and gesture did not function as hypothesized, possibly due to difficulty coding both these constraints.

By absolute value of coefficient, punctuation had the largest size of effect (-0.523). It was hypothesized that punctuation would have a clear and strong effect, as implied by the relevant section of Experiment 1; this is not a surprising result, though the fact that a significant result was achieved from so few stimuli is perhaps surprising in itself.

The second-largest coefficient by absolute value was lexical, with a coefficient of -0.284. The effect size is interesting, but the constraints lexical and phonotactic could not be analyzed as deeply as perhaps possible to those with more time and resources, as the constraints lexical and phonotactic are, by definition, entangled with the vowel and non-homerow vowel constraints. English phonotactics demands the presence of at least one vowel for lexical items. The relatively large size of Lexical’s coefficient, when analyzed outside of this possible convolution is, presents an interesting point of comparison for analysis that factors in these other constraints. Analysis that is able to properly account for these other constraints and their effects on each other may uncover a different picture of Lexical.

The length constraint had interesting relationships with constraints Keyset and Vowel. In the case of keyset, it was found that longer keysmashes may remain relatively acceptable while exhibiting some keyset violations; this was not so for the vowel constraint, where it was found that modeling violations remaining just as unacceptable regardless of keysmash length better fit the data. This cannot be due to a stronger effect of vowel (which had a coefficient of -0.143 compared to a coefficient of -0.212 for keyset). Properly investigating this interaction, and other interactions of length with other constraints, may be grounds for future work; understanding these interactions may also shed light on the nonlinear pattern of length acceptability ratings.

Repetition and Gesture had the murkiest results. Both these constraints were difficult to code and define, and oddities or inexplicabilities in the results may have resulted from these difficulties. In the case of Repetition, a variety of approaches to defining repetition were coded and analyzed. However, none of these constraints functioned quite as predicted. Most of the definitions of repetition had very small (yet significant) effects. The count of large skip-grams, for example, should have been more relevant than the presence or absence of skip-grams if Repetition was sensitive to the repetition of actual content in keysmashes. It is entirely possible that the non-significant results of the modified skip-gram approaches were not due to oddities in coding, but to the fact that the constraint they attempted to describe does not actually exist.

Issues with Gesture could result from either coding issues or an inaccurate concept of the constraint to begin with. Gesture was coded in such a way as to quantify the balance of left and right characters used in the keysmash. Though Gesture was hypothesized to be a sensitivity to the motion used to type the keysmash, the method used to code Gesture did not capture whether the gesture used to type the keysmash actually involves alternation between the hands; this lack could well explain the oddly positive effect that violations of the Gesture constraint were revealed to have. As the Gesture constraint was hypothesized as a response to keysmasher intuitions that some keysmashes with no violations of the other hypothesized constraints are more acceptable than others, it is also possible that there is simply no constraint related to hand gestures governing keysmash well-formedness. Polling more keysmashers on several more keysmashes that do not violate any of the other hypothesized constraints may be a good place to begin future work.

5. Limitations and Future work

Experiment 1 was designed to be as brief and as informative as possible, because excessive length would discourage participants from completing the survey; since taking the survey was entirely voluntary and motivated only by participants' curiosity and willingness to help, consideration of participants' attention spans was necessary. As Experiment 1's survey gathered 1243 responses in the span of approximately seven months, the design was likely successful in taking care to restrict its length in favor of participant volume. Experiment 2 took somewhat more of a risk in asking participants to judge 47 keysmashes for acceptability; though many more stimuli could have been utilized or created, the uncompensated nature of the task meant the survey had to be carefully designed for participants' attention span and willingness, and thus sacrificed some investigative power for a greater sample size. In environments where the survey could be circulated as a compensated research task (e.g. participants taking a survey for course credit or other compensation), this may not be as strong of a concern.

With keysmashes behaving linguistically according to at least three criteria, and with keysmashes demonstrating analyzable standards of form, the particular meanings of keysmashes come to light as a possible area for future work. Given more work in establishing keysmash standards of form, the changes that reflect different meanings can be differentiated from those that render a keysmash somewhat unusual or completely unacceptable. What keysmashes themselves can mean will first have to be re-examined: responses to the "other" fields in various checkbox-style questions of Experiment 1's survey, and responses to the tweet where this survey was initially posted, mention that the emotional categories given for keysmash meanings seemed too strict, and that keysmashes may have meanings that are between the given emotional options. In particular, Twitter replies often highlighted keysmashes used for sarcasm rather than to express a feeling. Keysmashes may used to convey a user's tone, attitude, or mood rather than a particular emotion they feel; investigating these uses and specific meanings may be fruitful.

In less farsighted matters, the single most pressing or relevant change to be made in future work concerns the design of stimuli. In future studies. ideally where the length of the task and the attention span of the participant is less of a concern, stimuli should be designed to provide more information about a single constraint, rather than designed to be compared within a transformation set, as in Experiment 2; or designed to cover several broad strokes of keysmash acceptability, as in Experiment 1. The transformation set design of Experiment 2's stimuli facilitated the comparison of various constraints to other constraints, but not so much the investigation of specific constraint effects. Most stimuli incurred a variety of constraint violations, rather than only violating one constraint per stimulus, which would have made a better basis for comparison.

Every constraint covered in Experiment 2 could easily be examined in more detail. The constraint Length is particularly interesting for the interactions it could potentially have with every other constraint; generally examining whether keysmashes remain acceptable with violations of a given constraint as the keysmash grows longer is a point of interest, as it is much more common to see numbers or special characters in keysmashes when those keysmashes are very long.

The Keyset constraint has a variety of approaches to further, detailed investigation, mostly pertaining to hypotheses examining the role of touch typists' hand positioning in acceptability ratings. If keysmash acceptability is strongly related to touch typing, keys from particular areas of the keyboard may be acceptable than others, if those areas are relatively more difficult to type. Similarly, this hypothesis would predict that keysmashes typed from mobile devices may have detectable differences when compared to those typed from keyboards, hinging on whether the majority of keysmashers position their hands on digital keyboards similarly to physical ones.

Hypotheses about different areas of the keyboard being more or less acceptable may have to account for the presence of vowels in various areas of the keyboard. The presence of three vowels at the top right of the keyboard in a standard QWERTY layout, combined with the relative unacceptability of vowels, may have an influence on keyboard area acceptability.

Though punctuation was shown to have a clear and negative effect on keysmash acceptability, other types of non-letter characters were not analyzed. Examining whether there is an effect particular to numbers or special characters, then comparing these potential effects, may be of interest. Specifically, if the inclusion of punctuation is less acceptable than the inclusion of numbers, there may be motivation to hypothesize that the evaluation of keysmashes by keysmashers may have to do more with the visual quality of the keysmash rather than the actual content of the keysmash itself. More obvious "interruptions" in the keysmash may be less acceptable than those less immediately visible (e.g. a potential difference between "SCL,JZF.S" and "SCL4JZF3S"). If this is the case, it may also explain the results obtained for the repetition constraint in this study.

Repetition may be a constraint sensitive to whether a keysmash *appears* to contain substrings that repeat. This could account for the fact that the repetition-boolean definition of repetition was the relatively best model among the five analyzed; this definition of repetition is not concerned with how many instances of repetition there are nor the size of the repeating skip-grams, only whether repeating skip-grams are present or not. The visual appearance of repetition, rather than the actual content of repetition,

An algorithm that is able to differentiate keysmashes from other words, and using this algorithm to process data scraped Twitter or other public social media contexts, would provide a good basis for a corpus analysis of keysmashes. This project could examine keysmash meanings in context for potential morphological analysis, or gather instances of attested keysmashes for further examination of constraints.

6. Conclusion

While keysmashes may appear random, Experiment 1 provides evidence that keysmashes instead convey meaning, have standards of form, have arbitrary relationships between those forms and meanings, and involve social conventions for their use. These behaviors align with criteria for language, and characterize keysmashes as linguistic phenomena. Experiment 2 provides evidence for factors governing keysmash well-formedness. After analyzing the results of a survey primarily focused on acceptability judgements, the hypothesized constraints vowel, length, keyset, punctuation, lexical and phonotactic proved to have statistically significant effects on keysmash acceptability ratings. In summary, keysmashes convey emotions and conversational tone that can be difficult to indicate over text-based communication, and despite their seemingly novel forms, keysmashes still behave linguistically. Even as language changes and evolves into unfamiliar forms, familiar linguistic tools can still be utilized to analyze and understand them as language. Keysmashes also challenge assumptions about what kinds of forms can convey linguistic meaning: though they can be dismissed as random-seeming utterances, the communicative value of keysmashes is more than conventional agreement about the meaning of a block of 'random characters'. The fact that keysmashers intuitively understand the well-formedness of keysmashes and distinguish various meanings between different instances show that this novel phenomenon behaves in a manner much more like language than it might otherwise seem. Thus, keysmashes, and perhaps other emergent phenomena of Internet communication, display a richness and complexity of linguistic expression that linguists have the tools to explore.

References

Bolte, C. (2021). *Ten-finger typing* [Digital image]. TypingAcademy. https://www.typing.academy/10-finger-typing

Hockett, C. F. (1960). The Origin of Speech. *Scientific American, 203*(3), 89-96. doi:10.1038/scientificamerican0960-88

Keysmash. (2020, November 10). Retrieved November 14, 2020, from https://en.wikipedia.org/wiki/Keysmash.

McCulloch, Gretchen. *Because Internet: Understanding the New Rules of Language*. Riverhead Books, 2019. (Keysmash practices, pgs. 6-7.)

Ygartua, H. (2018). Linguistic feature spread in online social networks. Undergraduate thesis at U of Oregon.

https://scholarsbank.uoregon.edu/xmlui/bitstream/handle/1794/24149/Final%20Thesis-Ygartua.pdf?sequence=1&isAllowed=y