SO INNOBET

Stylized Innovation in Synthetic Dictionaries

http://www.kinsler.org/physics/

Stylized innovation: interrogating incrementally available randomised dictionaries

Paul Kinsler*

Physics Department, Lancaster University, Lancaster LA1 4YB, United Kingdom. (Dated: Friday 15th June, 2018)

Inspired by recent work of Fink, Reeves, Palma and Farr [1] on innovation in language, gastronomy, and technology, I study how new symbol discovery manifests itself in terms of additional "word" vocabulary being available from dictionaries generated from a finite number of symbols. Several distinct dictionary generation models are investigated using numerical simulation, with emphasis on the scaling of knowledge as dictionary generators and parameters are varied, and the role of which order the symbols are discovered in.

PACS numbers:

I. INTRODUCTION

Fink et al. [1] recently presented an interesting discussion and analysis of serendipity and strategy in innovation scenarios, paying particular attention to real-world datasets based on language, gastronomy, and technology. Here, however, I take a look at the same type of innovation/discovery processes in the context of synthetic datasets where I have designed a range of different underlying "innovation" spaces to investigate. Although this is clearly not as realistic as using real-world datasets, it allows us to (a) try to establish both individual and universal features of a discovery/innovation process, and (b) avoids us having to make ad hoc judgments about whether chosen real items are individual components or compounds of other components.

In this paper I will investigate the synthetic datasets in a way broadly compatible with Fink et al. [1], but emphasizing discovery over determination of global properties, and relying primarily on numerical calculations. Notably, I will use some specific terminology that differs from theirs; this is in order to avoid unwanted inferences by clarifying the basic abstractions underpinning the dataset generation and its analysis. The total dataset will be called a *dictionary*, or world-dictionary, which is a list compiled of *words*. Each word in the dictionary is a string of *symbols* that has been randomly generated by an algorithm. As we *discover* extra symbols from the allowed list, we can *interrogate* the dictionary" of known words.

We will see how the process of ever-better interrogation of the dictionaries plays out in average terms and in some specific examples. Two different symbols discovery mechanisms are of interest here – most common first (i.e. frequency-order), and randomised. Of course, in many contexts we might think it most likely that the next-discovered symbol will be the most frequently occurring. Alternatively, when chasing innovation based on assembling known components, as might be true of cooking recipes or software systems, we can define an average usefulness [1] on which to base an intelligent symbolacquisition strategy.

However, when interrogating unknown systems, we may need to just rely for innovation on whatever new thing we have managed to discover. This will be without knowing in advance of use how frequently it will turn out to be useful – because we do not have access to the world-dictionary of our unknown system. This "explorative" view of innovation is distinct from that of the more "assembly"-focused view of innovation (combining known things) described above. Therefore we also show results for discovering the symbols in a random-order, where indeed we can see the sudden bursts of discovery we might normally expect from terms like "innovation" and "serendipity".

Unlike the treatment in Sood et al [2], these world dictionaries are static, and do not involve components (symbols) which appear as new, and then later may die off and disappear, taking the words they appear in with them. However, one might – although I do not do so here – consider interrogations where symbols are not only learnt but forgotten, so that the currently known sub-dictionary of the world-dictionary will never (or is unlikely to) reach a state of total knowledge. Instead, the subset of known symbols would only ever provide a sort of moving window on to the world-dictionary.

Other distinct models worthy of note here are the agentdecision led, persuasion/contagion-like models of Weiss et al [3], and the component-optimization view taken by McNerney et al [4] which led to nice power-law behaviour. However, as implied above, my model here follows more along the lines of "the system (dictionary) will only let me create a finite number of compounds (words), but I don't know what the allowed (symbols and words) are, and will only be able to find out as I discover new things (symbols)".

In what follows I will present and discuss the result of some numerical experiments based on algorithmically constructed dictionaries. Although real-world "dictionaries" are available, their underlying structure – the equivalent of the generating algorithm I use here to create my synthetic dictionaries – is unknown (or may not exist in any meaningful sense), which hampers analysis. In Sec. II I will present the generating algorithms along with the averaged scaling behaviour of discovery processes. However, since and real "innovation" process will only experience one possible symbol discovery order, in Sec. III I present and comment on specific discovery processes, noting their main features and discussing and comparing the outcomes. Finally, in Sec. IV, I present my conclusions.

II. DICTIONARIES AND SCALING

Here the dictionaries I investigate are compiled of words made from strings of symbols generated by a randomly driven algorithm. However, we might consider the symbols to repre-

•

^{*}Electronic address: Dr.Paul.Kinsler@physics.org

Stylized Innovation in Synthetic Dictionaries

http://www.kinsler.org/physics/

DI.Paul. KIIISIEI @ DIIVSICS.OIg

sent ingredients to be assembled into recipes ("words"), and then the compendium of recipes make up a recipe-book ("dictionary"); or they might be gestures and movements which we assemble into dance moves, and then into a whole dance routine. At any point in the discovery process, we will know only a subset of all the allowed symbols, and hence only have found ("innovated") a subset of the words; i.e. we will be working from not the total dictionary, but a sub-dictionary.

Here I consider how new words ("innovations") become available as new symbols are found and used to to interrogate the dictionary, and thus allow us to expand our sub-dictionary. For this we rely on aggregate quantities made up from all changes occurring during the process of discovering all symbols and hence the whole dictionary.

Fink et al [1] used an elegant series of diagrams that ranked symbols (their 'components') based on their "average usefulness $\bar{u}_{\alpha}(n)$ ", where the (ordinary) usefulness u_{α} of symbol α was a simple numerical count of how many words (their 'products') was present in. Here I prefer to focus on the usefulness as might be estimated due to the current state of knowledge, rather than an average usefulness [1] which involves a sum over all the alternative selections of other symbols. This less comprehensive averaging means we expect to see less stable rankings as the discovery process proceeds than Fink et al would, but it does not presume a knowledge of the (world) dictionary. Further, here I am most interested in the variation in the rates of change, not any absolute value.

The measures I consider when investigating the scaling of innovation are all normalized so that according to some simple assumptions, the full symbol discovery process on some dictionary will return that innovation measure to depend on the size of its symbol list, i.e. be N - 1. However, since the simple assumptions are in fact unlikely to hold in any given case, the difference between the calculated value and the nominal 'assumed value' will provide information about the innovation process. The measures are:

- 1. *Changes in ranks* δ_r , a count of how many symbolusefulness rankings changed at every new discovery, scaled down by *n*, i.e. the number of known symbols *n*. This is chosen so that if every known symbol changes rank at every new discovery, for each of the N - 1 discoveries, then $\delta_r = N - 1$.
- 2. *Linear ranks shift* δ_{ω} , a sum of all the absolute values of changes in symbol-usefulness at every new discovery, scaled down by $n^2/2$. This is chosen so that if every known symbol changes usefulness by an average n/2 at every new discovery, then $\delta_{\omega} = N 1$.
- 3. *Nonlinear ranks shift* δ_{χ} , a sum of all the changes in symbol-usefulness *squared* at every new discovery, scaled by $n^3/4$. This is chosen so that if every known symbol changes usefulness by an average n/2 at every new discovery, then $\delta_{\chi} = N - 1$; with the squaring acting to de-emphasize the smaller changes.

I now summarize and describe the dictionary generation methods I tested, and how the aggregate quantities varied with symbol count N, dictionary size W, and any generator-specific parameters. Each result in this section is an average taken from an ensemble of at least 16 test dictionaries, with more test dictionaries generated if necessary so as to keep the relative standard deviation in the aggregate quantities below 5%.

In the rest of this section all data plotted is summed over the entire discovery process, from first symbol to last.

A. Fake random dictionaries: "null" or N-dictionaries

These "null" dictionary are not real dictionaries – no word list exists, just a nominal word count of W and a nominal number of symbols N. They are defined solely as an algorithm intended to test the normalization of the δ quantities, and to provide a baseline of comparison that is unaffected by any specifics of a dictionary generation algorithm. If such an Ndictionary is interrogated, it returns information that the dictionary size is scaled in proportion to the number of known symbols (i.e. that w = nW/N, and that the allowed symbols have randomly ordered frequencies. Note that any reinterrogation of an N-dictionary will return a different randomised set of frequencies.

Fig. 1 compares the innovation measures for two symbol discovery possibilities: (a) discovery in order of frequency (or usefulness) in the world dictionary, or, (d) discovery in random-order. The two sets of results here span a range of symbol list and dictionary sizes. We see that for randomly chosen discovery order, $\delta_r \sim N$ within random error, as expected. The values for δ_{ω} and δ_{χ} are somewhat suppressed below the *N* value, as expected since it is very unlikely that any two subsequent interrogations will be as different as assumed by the normalization. As should be expected, the two discovery orders give results that differ only due to sampling error, since the "frequency-order" is randomly re-set by the algorithm at each subdictionary interrogation.

B. Fixed length words: F-dictionaries

A fixed length word dictionary $F_{N,\Omega}^W$ containing W words can be straightforwardly generated using a symbol list of length N, with Ω symbols per word, each symbol being chosen randomly. Words can contain any number (up to Ω) of repeated symbols, and words may be repeated in the dictionary.

With this generator, we will want $W \ll N^{\Omega}$ so that the random sampling of all possible words plays a strong role when populating the dictionary.

Fig. 2 compares the innovation measures for two symbol discovery possibilities: (a) discovery the frequency-order of symbols in the world dictionary, or, (b) discovery in randomorder. The two sets of results here span a range of symbol list and dictionary sizes, but are constrained by a fixed word length $\Omega = 8$. Since the dictionary parameters do always not honour the sampling inequality above, particularly for the smallest symbol lists, we see a pronounced 'knee' in each of the δ measures below N = 4.

For large symbol lists and small dictionary sizes we see that the δ measures report higher innovation for frequency-order discovery, although the difference is not large. This occurs because as the most frequent symbols occurs in the largest fraction of words, and so for any comparable shift in rank or frequency, the measures undergo a lesser rescaling due to the normalization process, since *n* is smaller earlier in the discovery process. Stylized Innovation in Synthetic Dictionaries

-œ/ INNOBET

DI.Paul. NIIISIEI @ DIIVSICS.OIg http://www.kinsler.org/physics/





FIG. 1: Scaling of innovation in N_N^W null random dictionaries, with symbol number N varying on a log scale from 2 to 32, and dictionary size W on a log scale from 45 to 1024.

Here, frequency-order discovery (open mesh) and random-order discovery (colour) results for Changes in ranks δ_r , Linear shifts in ranks δ_{ω} , and Nonlinear shifts in ranks δ_{χ} . Since the open mesh is often hard to see easily, the colour-coding of the random-order result surface is done according to the difference between the two sets of data.

Perhaps most interestingly, we see that the change count index δ_r reports higher innovation for larger symbol lists, in a way approximately with its definition, but this behaviour is suppressed for δ_{ω} , and even reversed for the nonlinear measure δ_{γ} . This implies that the discovery process is more incremental for larger N, involving relatively few(er) discoveries that substantially revolutionize the size and structure of the currently known sub-dictionary. This is not unexpected, since in F-dictionaries with larger word counts, the whole set of symbols will be better and more evenly sampled.

Since the above results were calculate for a fixed word length, in fig. 3 I instead hold the dictionary size W fixed at 1024, and vary word length Ω . There is a clear trend for higher innovation for shorter words, and the trend for suppressed innovations δ_{ω} and δ_{χ} at larger N systematically fades away for short words, and even reverses for the shortest word lengths. This is most likely caused by the same reason there was a knee for small dictionary sizes in fig. 2: e.g. taking opposite corners of the parameter ranges, note that with N = 32 and $\Omega = 2$,

FIG. 2: Scaling of innovation in $F_{N,8}^W$ fixed random dictionaries, with symbol number N varying on a log scale from 2 to 32, and dictionary size W on a log scale from 45 to 1024.

Here, frequency-order discovery (open mesh) and random-order discovery (colour) results for Changes in ranks δ_r , Linear shifts in ranks δ_{ω} , and Nonlinear shifts in ranks δ_{χ} . Since the open mesh is often hard to see easily, the colour-coding of the random-order result surface is done according to the difference between the two sets of data.

or for N = 2 and $\Omega = 10$, the number of unique words is 32^2 or 2^{10} – both being the same as W = 1024.

C. Extensible: E-dictionaries

An extensible dictionary E_N^W containing W words is generated using a symbol list of length N. Each new word is added iteratively, starting from a single word consisting only of the initial symbol α_{\circ} , randomly chosen new symbols are appended until the result is not already in the dictionary. There is no maximum word length and no generator-specific parameters.

Fig. 4 compares the innovation measures for two symbol discovery possibilities: (a) discovery the frequency-order of symbols in the world dictionary, or, (d) discovery in randomorder. The two sets of results here span a range of symbol list and dictionary sizes. For low word counts, and large enough symbol lists, some symbols might not be chosen for use in

Stylized Innovation in Synthetic Dictionaries





FIG. 3: Scaling of innovation in $F_{N,\Omega}^{1024}$ fixed random dictionaries, with symbol number *N* varying on a log scale from 2 to 32, and word length Ω from 2 to 11.

Here, frequency-order discovery (open mesh) and random-order discovery (colour) results for Changes in ranks δ_r , Linear shifts in ranks δ_{ω} , and Nonlinear shifts in ranks δ_{χ} . Since the open mesh is often hard to see easily, the colour-coding of the random-order result surface is done according to the difference between the two sets of data.

FIG. 4: Scaling of innovation in E_N^W extensible dictionaries, with symbol number *N* varying on a log scale from 2 to 32, and dictionary size *W* on a log scale from 45 to 1024.

Here, frequency-order discovery (open mesh) and random-order discovery (colour) results for Changes in ranks δ_r , Linear shifts in ranks δ_{ω} , and Nonlinear shifts in ranks δ_{χ} . Since the open mesh is often hard to see easily, the colour-coding of the random-order result surface is done according to the difference between the two sets of data.

D. Chain: C-dictionaries

words in the dictionary; the average number of unused symbols is indicated by the floating contours on the figure.

For large symbol lists and small dictionary sizes we see that the δ innovation measures report higher innovation for frequency-order discovery, although the difference is not large. The likely reason is the same as for F-dictionaries, and due to the weaker effect of the normalization on early innovative bursts.

Despite the very different dictionary generation algorithm, the trends in δ_r are rather similar to those seen in Fdictionaries, although the measure is more strongly depressed for small dictionaries with large symbol lists. However, the linear and nonlinear measures δ_{ω} , δ_{χ} behave very differently, being in line with the 'expected' behaviour of increasing with symbol list size. A chain dictionary $C_{N,v}^W$ containing W words is generated using a symbol list of length N. It is an incremental process, based on a starting point consisting of an initial random symbol α_o . Each new word is added iteratively: either by (i) with probability 1 - v, an existing word being chosen and a new random symbol appended to create the new word, which is only kept if the result is not already present; or (ii) with probability v, another (single) random symbol is chosen to be the new word, which is only kept if the result is not already present. In this way the "fork probability" v determines how fast the dictionary accumulates words with different starting symbols. Although we should note that the starting symbol has no special role, a low v means that the early-chosen symbols, and in particular α_o , will dominate the symbol frequencies in the resulting dictionary.

For low word counts, and large enough symbol lists, some symbols might not be chosen for use in words in the dictio**TO** INNOBET

Stylized Innovation in Synthetic Dictionaries

http://www.kinsler.org/physics/

nary.





FIG. 5: Scaling of innovation in $C_{N,v}^W$ chain dictionaries, with symbol number *N* varying on a log scale from 2 to 32, and dictionary size *W* on a log scale from 45 to 1024.

Here, frequency-order discovery (open mesh) and random-order discovery (colour) results for Changes in ranks δ_r , Linear shifts in ranks δ_{ω} , and Nonlinear shifts in ranks δ_{χ} . Since the open mesh is often hard to see easily, the colour-coding of the random-order result surface is done according to the difference between the two sets of data. The floating contours that appear indicate the number of symbols that (on average) are not chosen to appear in the dictionaries.

Fig. 5 compares the innovation measures for two symbol discovery possibilities: (a) discovery the frequency-order of symbols in the world dictionary, or, (d) discovery in randomorder. The two sets of results here span a range of symbol list and dictionary sizes, but are constrained by a fixed word fork probability v = 10%. For low word counts, and large enough symbol lists, some symbols might not be chosen for use in words in the dictionary; the average number of unused symbols is indicated by the floating contours on the figure.

There is no strong trend with dictionary size, but in a new feature not present for the previous dictionary types, discovery order now plays a very strong role: discovery in frequency-order shows the expected innovation increase with N, whereas for random-order discovery the trend is weak (for δ_r), or negligible (for $\delta_{\omega}, \delta_{\chi}$). Once might, for example, even hazard a comparison between the behaviour of F-dictionary behaviour and this for the random-order discovery; despite the fact that

FIG. 6: Scaling of innovation in $C_{N,v}^{1024}$ chain dictionaries, with symbol number *N* varying on a log scale from 2 to 32, and fork probability *v* from 1/11 to 10/11.

Here, frequency-order discovery (open mesh) and random-order discovery (colour) results for Changes in ranks δ_r , Linear shifts in ranks δ_{ω} , and Nonlinear shifts in ranks δ_{χ} . Since the open mesh is often hard to see easily, the colour-coding of the random-order result surface is done according to the difference between the two sets of data.

the frequency-order discovery looks more like that for N- or E-dictionaries! Indeed, this is again probably a signature of the normalization effect: the first (or first few) symbols used in the dictionary generation process will appear in the largest fraction of words. In frequency order discovery these innovation bursts occur early and relatively isolated, in random order they may well appear late and amongst a significant background. Naturally, such effects stand out more in small dictionaries, where change plays a greater role. The strong discover-order effect is also seen on fig. 6, where *W* is held fixed at 1024, and fork probability *v* is varied: for *v*, the early generated symbols are more dominant, and so randomness in the generation process plays a stronger role.

E. Blinkered: B-dictionaries

A Blinkered dictionary $\mathsf{B}_{N,\nu}^W$ containing W words is generated using a symbol list of length N. It is an incremental pro-

Stylized Innovation in Synthetic Dictionaries

http://www.kinsler.org/physics/

cess, based on a starting point consisting of an initial random symbol α_{\circ} , and is somewhat similar to a chain dictionary, but with new symbols only arising due to forks. Each new word is added iteratively: either by (i) with probability 1 - v, one existing word is randomly chosen, a second randomly chosen existing word is appended; but is only kept if the result is not already present; or (ii) with probability v, another (single) random symbol is chosen to be the new word, but is only kept if the result is not already present.

In this way the fork-fraction v determines how fast the dictionary accumulates words with different symbols. Although we should note that the starting symbol has no special role, a low v means that the early-chosen symbols, and in particular α_{\circ} , will dominate the symbol frequencies in the resulting dictionary.



FIG. 7: Scaling of innovation in $B_{N,v}^W$ blinkered dictionaries, with symbol number *N* varying on a log scale from 2 to 32, and dictionary size *W* on a log scale from 45 to 1024.

Here, frequency-order discovery (open mesh) and random-order discovery (colour) results for Changes in ranks δ_r , Linear shifts in ranks δ_{ω} , and Nonlinear shifts in ranks δ_{χ} . Since the open mesh is often hard to see easily, the colour-coding of the random-order result surface is done according to the difference between the two sets of data. The floating contours that appear indicate the number of symbols that (on average) are not chosen to appear in the dictionaries.

Fig. 7 compares the innovation measures for two symbol discovery possibilities: (a) discovery the frequency-order of symbols in the world dictionary, or, (d) discovery in random-



FIG. 8: Scaling of innovation in $B_{N,v}^W$ blinkered dictionaries, with symbol number *N* varying on a log scale from 2 to 32, and fork probability *v* from 1/11 to 10/11.

Here, frequency-order discovery (open mesh) and random-order discovery (colour) results for Changes in ranks δ_r , Linear shifts in ranks δ_{ω} , and Nonlinear shifts in ranks δ_{χ} . Since the open mesh is often hard to see easily, the colour-coding of the random-order result surface is done according to the difference between the two sets of data.

order. The two sets of results here span a range of symbol list and dictionary sizes, but are constrained by a fixed word fork probability v = 10%. For low word counts, and large enough symbol lists, some symbols might not be chosen for use in words in the dictionary; the average number of unused symbols is indicated by the floating contours on the figure. Note that the number of unused symbols tends to be very much higher than for the E- or C- dictionaries, this is because a successful "fork" is now the *only* way for new symbols to appear in the evolving dictionary. Nevertheless, despite this distinct unused-symbol behaviour, the average number of symbols extant in the dictionaries still increases with increasing N.

For these B-dictionaries, there is no strong trend with dictionary size, unlike for F- and E-dictionaries; but now higher innovation δ_r is reported for random discovery order, in this is is like the N-dictionary results; but this is not seen for δ_{ω} or δ_{χ} .

To address the effects of fork probability, in fig. 8 I instead hold W fixed at 1024, and vary v. The main feature of note

- MINNOBET

Stylized Innovation in Synthetic Dictionaries

http://www.kinsler.org/physics/

here is that there is a downward trend in reported innovation δ_r for small ν , especially for large symbol lists; in this it is similar to a C-dictionary.

F. Discussion

An important feature of the results shown here is the role of discovery order, and how for some dictionary generation algorithms it plays a key role, but for others it does not. In particular, the significance of discovery order for C-dictionaries is rather large, but for B-dictionaries less so – this being despite the fact that the generation of B-dictionaries is much more contingent on the generation order of symbols, and so we might expect that to influence the effect of discovery order also. Of course, a strong caveat here is that these are results *averaged* over many dictionaries, and so do not give a clear picture of the experience of any one specific symbol discovery trajectory. Consequently, I address this point in particular in the next section.

III. DISCOVERY

In the previous section I looked at the measures if innovation, as averaged over the entire discovery process from beginning to end. However, as stated in the introduction, here I want to investigate innovation as a dynamic process evolving into an unknown future. So, with the aggregated results of the previous section as a backdrop, I will now look at some specific examples, based on a given dictionary but different discovery sequences. This will usually be a frequency-order discovery of new symbols, as well as two different randomorder sequences. There are of course other options – the pessimistic might try reverse-frequency-order, where little-used symbols are discovered first; or a frequency-weighted random discovery process.

The data plotted are as follows:

- (a) Change in average symbol frequencies: the average and standard deviation of the symbol frequencies at each discovery event are calculated, then used to display the *change* in these between such knowledge at this event and the previous event. This is plotted on a log scale, with 1 added to the absolute value of the change, so that any visible peaks are indicative of a significant innovation event where many new words become available. A higher curve, being the average value plus one standard deviation (in the mean) is also shown.
- (b) Symbol entropy, calculated as $E_n = -\sum p_{i,n} \log_2 p_{i,n}$, where $p_{i,n}$ is the probability of symbol *i* being in a word knowable at the *n*-th symbol discovery.
- (c) Symbol rankings: However, these are not the raw rankings according to the latest information from only this discovery event, but are averaged over the process so far, then re-ranked. Processing is done according to the default method in R, so that tied rankings are given an average value (i.e. a tie at 3 and 4 results in both being ranked 3.5).

A. Fixed length words

Here we discuss the symbol ranking evolution and known word counts for a sample $F_{32,8}^{1024}$ F-dictionary. Since the generation algorithm and word length chosen means that it is extremely unlikely that a word consisting of a single symbol will be in the dictionary, we see that many symbol discoveries are needed before *any* words from the dictionary are known. This delayed onset is a particular characteristic property of this kind of fixed-length word dictionary – or, indeed, of related types of dictionary which contain only finite minimum length words of randomly chosen symbols. Another characteristic is that – once started – the amount of word innovation accelerates as more symbols are discovered. However, note that this acceleration in innovation is reliant on the size of the dictionary being much less than the number of possible words.

The abbreviated symbol entropy curves occur because it cannot be calculated for sub-dictionaries which contain no words.

1. Frequency order discovery (F)

For frequency-order symbol discovery, the main point to note is that the discovery process is highly dynamic, with frequent changes in symbol ranking. However, these many changes average out to provide relatively smooth progress in the resulting word discovery rate. Since the most frequent symbols are in so many words, the start of word innovation is delayed. In the example shown in fig. 9, it is delayed until 8+ symbols are known – and there is another lesser holdup until about the 15th discovery. These bottlenecks are visible not only in the word rankings, but also the lack of changes in average frequency, and by the premature plateau in symbol entropy.

2. Random discovery (F)

In some respects, random discovery provides a rather similar picture to the frequency-order symbol discovery – after a sudden initial onset, there are many ranking changes, but a smooth progress in the resulting word innovation. This is unsurprising, since new symbols were added randomly to the dictionary as it was grown and so variations in symbol frequency are generally small, as can be seen from the very small standard deviation.

By chance, the two random discovery process shown in fig. 10 have given distinct outcomes – one has the double innovation burst also seen in frequency-order discovery, whereas the other does not. Nevertheless, both examples have an even more delayed onset than frequency order discovery, being later than the tenth discovery; such behaviour is caused by initial discoveries being of symbols that do not appear together. This behaviour is not unexpected, since the number of possible words is very much greater than the chosen size of the dictionary.

DI.Paul. NIIISIEI @ DIIVSICS.OIg

SO INNOBET

Stylized Innovation in Synthetic Dictionaries

rd innovation increases much more slowly as



FIG. 9: Summary of innovations from a sample $F_{32,8}^{1024}$ fixed Fdictionary, given frequency-order discovery. Upper frame: The black line shows the change in average symbol frequency at each new discovery, with the dotted line showing that average plus one standard deviation in the mean. The red dashed line shows the symbol entropy E_n , with the red bar on the right indicating the range from zero to $\log_2(32)$. The green dot-dashed line shows the fraction of words discovered. Lower frame: The coloured lines show the symbol frequency rankings, as averaged over previous discovery steps, then reranked. Identical ranks are averaged, so that (e.g.) if two symbols are ranked equal third, the average rank is 3.5.



FIG. 10: Summary of innovations from the sample $F_{32.8}^{1024}$ fixed Fdictionary. Results for two different randomised symbol discovery orders are shown. Although individual details differ, the data shown on the upper frames of these plots have a character similar to the frequency-order discovery. In contrast, as can be seen in the lower frames, there are very many more changes in symbol ranking, How the data is plotted is described in the caption to fig. 9.

B. Extensible

Here we discuss the symbol ranking evolution and known word counts for a sample E_{32}^{1024} E-dictionary. Since the generation algorithm means that single symbol words exist, word innovation begins immediately, unlike for the delayed start seen in the F-dictionary case. Another characteristic the

amount of word innovation increases much more slowly as more symbols are discovered – at least as compared to the previous F-dictionary case.

1. Frequency order discovery (E)

For frequency-order symbol discovery, the main point to note is that the discovery process shown on fig. 11 starts immediately, and continues with frequent changes in symbol ranking. However, as for F-dictionaries, these many changes average out to provide gradual progress in the resulting word innovation rate. While we see that indeed some "innovation" occurs, as reflected in the swapping of symbol rankings in the lower frame plot, the word increase rate on the upper frame shows no sudden increases in discovered-word counts. Note also that some of the most frequent symbols do not stay as the most highly ranked for the whole discovery process. Further, the symbol entropy increases smoothly towards its maximum value throughout the discovery process.



FIG. 11: Summary of innovations from a sample E_{32}^{1024} extensible Edictionary, given frequency-order discovery. How the data is plotted is described in the caption to fig. 9.

2. Random discovery (E)

Example of random discovery, as shown on fig. 12 provide a rather similar picture to the frequency-order symbol discovery. This is unsurprising, since new symbols were added randomly to the dictionary as it was grown and so variations in symbol frequency are generally small, as can be seen from the small standard deviation. Nevertheless, we do see that the number of ranking changes is increased in comparison to frequency order discovery, although there is little sign of this dynamics in the frequency changes or symbol entropy.

Stylized Innovation in Synthetic Dictionaries

http://www.kinsler.org/physics/



FIG. 12: Different innovation histories from a single sample E_{32}^{1024} extensible E-dictionary. Results for two different randomised symbol discovery orders are shown, each with their own distinct experiences of word innovation. Although individual details differ, the data shown on the upper frames of these plots have a character similar to the frequency-order discovery, especially for the symbol entropy. In contrast, as can be seen in the lower frames, there are very many more changes in symbol ranking, How the data is plotted is described in the caption to fig. 9.

C. Chain

Here we discuss the symbol ranking evolution and known word counts for a sample $C_{32,0.1}^{1024}$ chain C-dictionary. Just as for the E-dictionary, since the generation algorithm means that single symbol words exist, word innovation can easily begin immediately, unlike for the F-dictionary case.

We see here for the first time a clear difference between the frequency-order and random-order discovery processes. Both have bursts of innovation, but since the random-order discovery process is unlikely to begin with the symbols used early on in the generation process, we have the potential for a delayed start to word innovation while we wait for sufficiently common symbols to be found. Further, the bursts of innovation are stronger in random discovery, not only adding more structure to changes in symbol rankings, but even slowing or interrupting the increase in symbol entorpy.

1. Frequency order discovery (C)

Here, new symbols are discovered by interrogators in order of their frequency of occurrence in the total dictionary. While we see on fig. 13 that indeed some "innovation" occurs, as reflected in the swapping of symbol rankings in the lower plot, the upper plot does show some variation – perhaps best described as brief bottlenecks – in the increase in symbol frequencies. However, we can see that whilst the symbol rankings remain more stable here that for either the F- or Edictionaries, the changes in discovered-word counts are more variable.

2. Random discovery (C)

Unlike the frequency-order discovery for these Cdictionaries, and for the other F- and E-dictionaries, here in fig. 14 we do see sudden bursts of word innovation, with strong variation in changes in symbol frequency. Further,



FIG. 13: Summary of innovations from a sample $C_{32,0.1}^{1024}$ chain C-dictionary, given frequency-order discovery. How the data is plotted is described in the caption to fig. 9.

although the symbol entropy does not exhibit the step-like change that can be seen for F-dictionary results, it is possible that discovery of a new symbol (α_{\star}) permits so many new words to be found (i.e. colloquially "unleashes a wave of innovation") that its ubiquity causes a *drop* in the symbol entropy, as is marked by the vertical bar on fig. 14(right). This phenomenon is prefigured by a significant period where new symbol discoveries have little effect on the symbol frequencies – presumably these at-the-time apparently uninteresting symbols mostly appear in words that also need the thenunknown α_{\star} .



FIG. 14: Different innovation histories from a single sample $C_{32,0.1}^{1024}$ chain C-dictionary, the same one as for fig. 13 above. Results for two different random symbol discovery orders are shown, each with their own distinct experiences of word innovation. The data shown on the upper frames of these plots have a much more jagged character than to the frequency-order discovery, and on the (right) 0099 plot the first strong innovation burst is even associated with a small drop in symbol entropy (as marked by a vertical red line). In common with result from previous dictionary types, there are very many more changes in symbol ranking, as compared with frequency-order discovery. How the data is plotted is described in the caption to fig. 9.

Stylized Innovation in Synthetic Dictionaries

D. Blinkered

Here we discuss the symbol ranking evolution and known word counts for a sample $B_{32,0.2}^{1024}$ blinkered B-dictionary. Not unexpectedly, if we consider the dictionary generation process, the distinction between frequency-order and random-order discovery is by far the clearest here.

When new symbols are discovered by interrogators in frequency-order, innovation increases rapidly at the start but then flattens out. In contrast, even with eight possible randomorder discovery processes to choose from, there seems to be no commonality between the histories. There may or may not be a delayed start to innovation, there can be isolated bursts of innovation followed by stagnation, there can be multiple nearby bursts, and these might or might not take place on a background level of innovation.

Lastly, another new feature seen in B-dictionaries is that innovation bursts can cause the symbol entropy E_n to under go a small decrease, which briefly counteracts its usual steady increase. For all other dictionaries, the increase in E_n was monotonic.

1. Frequency order discovery (B)

We can see in fig. 15 a very distinct behaviour – most notably, very few changes in ranking. This is to be expected, given that the inclusion of new symbols was entirely controlled by the fork probability, and the discovery order. A consequence of this is that the change in discovered-word counts tails off significantly as the last few – and rare – symbols are discovered. In fact, because of the way these Bdictionaries are created, frequency-order discovery will match quite closely to the order of appearance of new symbols as the dictionary was generated – especially for the beginning of the discovery process.

2. Random discovery (B)

Unlike the frequency-order discovery, here we again see sudden bursts of word innovation, as for the random-discovery C-dictionaries, and are shown on fig. 16. However, as the ranking swaps occur, they do so in a more structured way, and many new words become available. These bursts of word innovation will largely be due to the discovery of symbols used early in the dictionary generation process, which we would expect to appear in a large fraction of possible words. Indeed, the dictionary generation process is such that we tend to see rather more of the "wave of innovation" induce symbol entropy drops in the discovery history than for C-dictionaries. Again, each is generally prefigured by a period where new symbol discoveries seem to have little effect, even though in fact each of those apparently unremarkable discoveries act to prime our knowledge base for the burst of innovation made possible for when the key symbol is discovered.

E. Discussion

One point to note is that when comparing different measures of the discovery process, we that what might look like



FIG. 15: Summary of innovations from a sample $B_{32,0.2}^{1024}$ blinkered Bdictionary, given frequency-order discovery. How the data is plotted is described in the caption to fig. 9.

innovation under one measure may not necessarily seem like it under another. In the work of Fink et al [1] it was changes in rank that was emphasised as a signature of innovation, but it is arguable that other measures – perhaps peaks in the change in symbol frequencies, or the brief dips in the symbol entropy provide a better guide. To determine which is appropriate in any given case needs some thought about what we mean by innovation. This may, of course, be influenced by the underlying dictionary generation model - e.g. in F-dictionaries, each of any step increases in the symbol entropy looks like a promising signature ... but in B-dictionaries, it is the brief *dips* that appear to be the key. Still, one advantage of symbol entropy is that either signature is either present or not – whereas the perhaps more robust indicator, i.e. large changes in symbol frequencies, requires some ad hoc definition of the threshold that needs to be passed to be called a "burst of innovation".

A feature of the discovery process which may help determine the appropriate measure is the discovery order. Here I have used frequency-order or random-order, which has been shown to strongly influence some proposed innovation signatures. However, real discoveries are unlikely to be made in the (eventual) frequency order, but it seems implausible that they would be truly random either; although this judgement would likely depend on what sort of innovation space was under discussion. Perhaps with software component "symbols", frequency-order might be reasonable, if you assume that effort is spent on solving common problems first, so that those components end up being written ("discovered") - more or less according to how much they are eventually used. In contrast, if applied to a model of basic research, some common problems ("symbols") might be so hard to solve that regardless of effort, other much easier problems are discovered first.

-Ø∣

Stylized Innovation in Synthetic Dictionaries

http://www.kinsler.org/physics/



FIG. 16: Different innovation histories from a single sample $B_{32,0.2}^{1024}$ blinkered B-dictionary, the same one as for fig. 15 above. Results for eight different random symbol discovery orders are shown, each with it's own distinctive pattern of average frequency changes and symbol ranking shifts. Further, drops in symbol entropy here are both more common than for other dictionaries, and can have larger values. How the data is plotted is described in the caption to fig. 9.

IV. CONCLUSION

We have seen that the effect of symbol discovery allows new word innovation in all the model dictionary generators used here. However, the way the dictionary is generated alters the nature of innovation. The fixed random F-dictionaries required many symbols to be researched before any words could be found, after which word innovation improved gradually with each new symbol. Further, although the mechanism of symbol discovery does affect that process in individual cases, it is not a strong effect on average or in general. The extensible E-dictionaries are more forgiving, with their inclusion of short words allowing early word innovation, but otherwise innovation tended to be gradual, and not on average particularly affected by the mechanism of symbol discovery.

In contrast, the chain C-dictionaries can allow the expectation of bursts of innovation, especially under the random symbol discovery model. This is seen on average, in the scaling data of section II and also in the sample data of section III. Notably, we can see such innovation bursts when key symbols are discovered in fig. 14. This characteristic behaviour is cause by its generation process, which – especially for low fork probabilities v – tends to reuse existing words, until a forking event allows new possibilities to emerge. The crucial event in the discovery process, and then for the number of subsequently available innovations, is therefore finding the symbol which started the generation algorithm.

This phenomenon is even more pronounced in blinkered Bdictionaries, which are even more dominated by the first (and other early) symbols. Here, the number of available innovations remains restricted until those particular symbols are discovered.

This paper has attempted to evaluate how the nature of the underlying innovation space (here, *dictionary*) might reveal itself as incremental discoveries (i.e. of components or *symbols*) are made, and new innovations (*words*) become available. The variation in results has shown that there is indeed some signature, but whether we could expect it to be visible above the variations caused by the ad hoc discovery of new components in our single instance of a discovery process is debatable. In particular there is no suggested quantitative measure one might use to try to infer some properties of the underlying innovation space; although the results – to an extent – might provide some small basis for generating some intuition.

To summarize, this work leaves us with two tricky problems to address – 'Which innovation measure is best?', and 'What discovery order is most appropriate?'; to which we should also add a third – 'Given some partial discovery/innovation history, is it possible to infer the nature of the underlying dictionary?'. An answer to this last question – assuming it were possible to find one – might perhaps enable us to optimize our ongoing discovery effort. All these three questions however are ones that apply to real world discovery/innovation efforts, and not the stylized ones presented here; nevertheless, this purely synthetic approach had the advantage of being fully transparent to investigation.

Acknowledgements

I am grateful for the support provided by EPSRC (the Alpha-X project EP/N028694/1).

Stylized Innovation in Synthetic Dictionaries

http://www.kinsler.org/physics/

Methods

The dictionary generation, analysis, and plotting was done with scripts written for and run using Scilab [5] and R [6].

- T. M. A. Fink, M. Reeves, R. Palma, and R. S. Farr, Nat. Commun. 8, 2002 (2017), arXiv:1608.01900, doi:10.1038/s41467-017-02042-w.
 Y. Saad, M. Mathim, A. Shmim, B. Canada
- [2] V. Sood, M. Mathieu, A. Shreim, P. Grassberger, and M. Paczuski, Phys. Rev. Lett. 105, 178701 (2010),
- doi:10.1103/PhysRevLett.105.178701.[3] C. H. Weiss, J. Poncela-Casasnovas, J. I. Glaser, A. R. Pah, S. D.
- [5] C. H. Weiss, J. Fonceta-Casashovas, J. F. Giaser, A. K. Fan, S. D. Persell, D. W. Baker, R. G. Wunderink, and L. A. N. Amaral, Phys. Rev. X 4, 041008 (2014),

doi:10.1103/PhysRevX.4.041008.

- [4] J. McNerney, J. D. Farmer, S. Redner, and J. E. Trancik, PNAS **108**, 9008 (2011), doi:10.1073/pnas.1017298108.
- [5] *Scilab*, Scilab Consortium, software, Open Source, http://www.scilab.org/.
- [6] *R*, R Foundation, software, Open Source, https://www.r-project.org/.