SweetPea: A standard language for factorial experimental design.

Sebastian Musslick^{1,*},

Anastasia Cherkaev²,

Ben Draut²,

Ahsan Butt²,

Vivek Srikumar²,

Matthew Flatt²,

Jonathan D. Cohen^{1,3}

¹Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA.

²School of Computing, University of Utah, Salt Lake City, UT 84112, USA.

³Department of Psychology, Princeton University, Princeton, NJ 08544, USA.

*Correspondence: musslick@princeton.edu

Abstract

Experimental design is a key ingredient of reproducible empirical research. Yet, given the increasing complexity of experimental designs, researchers often struggle to implement ones that allow them to measure their variables of interest without confounds. SweetPea is an open-source declarative language in Python, in which researchers can describe their desired experiment as a set of factors and constraints. The language leverages advances in areas of computer science to sample experiment sequences in an unbiased way. In this article, we provide an overview of SweetPea’s capabilities, and demonstrate its application to the design of psychological experiments. Finally, we discuss current limitations of SweetPea, as well as potential applications to other domains of empirical research, such as neuroscience and machine learning.

Keywords: factorial design; randomization and sampling; sequential constraints; nuisance factor

SweetPea: A standard language for factorial experimental design.

Introduction

Reproducibility is an issue of increasing concern in scientific research. The problem has gained considerable attention in psychological research, where behavioral phenomena are often confounded by the presence of various moderator variables (Sherman & Pashler, 2019; Klein et al., 2014; Collaboration et al., 2015; Stroebe & Strack, 2014; Dijksterhuis, Van Knippenberg, & Holland, 2014). More generally, the problem arises in any field where an experiment systematically manipulates a set of variables, controlling as best as possible for variables of non-interest, to evaluate the effects of experimental manipulation. This approach, a fundamental staple of the scientific process, applies to empirical studies in the natural sciences as well as the training and evaluation of artificial systems in engineering, such as machine learning algorithms (Langley, 1988; Zadrozny, 2004; Drummond, 2006; Bergstra & Bengio, 2012; Gardner, Neumann, Grus, & Lourie, 2018).

Experimental Design

Methods of experimental design date back to the study of inheritance through the cultivation of pea plants¹ (Miller, 2011), and remains a gold standard in psychological research. In experimental settings, an ideal design samples variables of experimental interest in an unbiased way, while suitably controlling for variables of non-interest. While this is easy to state in principle, it is often difficult to achieve in practice. This can be for many reasons. One is that it may be impossible to sample all of the combinations of variables needed to achieve a fully balanced design. Random sampling is often used to address this. However, when data collection is expensive and therefore sample sizes are constrained, random sampling may be inadequate to ensure an adequately balanced or controlled design (Button et al., 2013; Kühberger, Fritz, & Scherndl, 2014; Rossi, 1990; Wells & Windschitl, 1999). Randomization is also often

¹ This involved the crossing of pure-breeding lines of plants as experimental factors and assessing traits of the hybrid progeny (Druery & Bateson, 1901; Mendel, 1866).

used when the complexity of a design makes it difficult to determine how to satisfy constraints imposed by the design in a more systematic way (e.g. Mayr & Keele, 2000); however, naive randomization can render the design vulnerable to unintended and unrecognized biases. There is an increasing risk that this problem will be amplified by platforms that provide the infrastructure for automated data collection (De Leeuw, 2015; Gureckis et al., 2016; Hartshorne, de Leeuw, Goodman, Jennings, & O'Donnell, 2019); without proper experimental design, these approaches to large-scale data collection may encounter the same issues of reproducibility as traditional data collection, but in more arcane and difficult to detect ways. Finally, interest is increasing in optimal experimental design where, rather than simple counterbalancing, it is helpful to identify experimental conditions that are best for discriminating between models (Dale, 1999; Myung & Pitt, 2009)

These issues also arise in engineering. For example, in machine learning, the outcome of training is profoundly influenced by the statistics of the stimuli on which the model is trained; small sample sizes or violations of proper random sampling can limit performance or validity of statistical learning techniques, such as naive Bayes or decision tree learners (Zadrozny, 2004) or, conversely, it may be useful to identify designs that contrast conditions or segment the data in ways that optimize learning. Furthermore, the success of a machine learning system crucially depends on the choice of feature representations, models and their hyperparameters, which together call for searching through a potentially intractable search space (Chapelle, Vapnik, Bousquet, & Mukherjee, 2002; Bergstra & Bengio, 2012).

Factorial Structure vs. Implementation Constraints

The best approach to the concerns raised above is a well designed experiment. Factorial structure is the cornerstone of well-controlled experimental design involving repeated measures: The experimental variables manipulated in each repeated measure, or *trial*, are defined as *factors*, and the values of each variable to be sampled in the experiment are defined as *levels* of the corresponding factors. The trials of the

experiment are then constructed by crossing the factors – that is, selecting the levels from each factor to be used in each trial. In addition to providing a concise, formally rigorous description of the experimental design, expression in factorial form also aligns with standard approaches to statistical analysis; for example, each factor can be represented as an effect in an analysis of variance (ANOVA) or, more generally, as a regressor in an analysis using the general linear model (GLM).

When an experimental design consists of a small number of independent factors, each with a modest number of levels, and all combinations of their levels can be subjected to measurement, then generating the trials for an experiment is straightforward: cross all of the factors, generate every possible permutation of the crossings (i.e., sequences of trials), and use each permutation in a different run of the experiment (this is often referred to as *full counterbalancing*). However, for most designs this is unrealistic for one or more of several reasons. First, the full crossing of all factors may generate too many combinations of levels to sample in a single run of the experiment (e.g., for a single participant to complete). Second, even if the full crossing can be sampled in a single run, there may be too many permutations to sample over runs of the experiment (e.g., too many participants would have to be run). These challenges are often met by writing a program that sub-samples from the full crossing and permutation of sequences. Thus, for the first problem, the program might select a subset of combinations of levels from the full crossing for use in each run, either randomly or by insuring that all possible subsets are used across runs of the experiment (e.g., using a Latin Squares design). For the second problem, sequences can be drawn randomly from the full set of permutations (of either the full crossing, or the subset of it used for each run). Sometimes such procedures are sufficient to achieve an adequately balanced sample of the different implementations of the design, in which case measurements across runs are not biased in way(s) that confound interpretation. However, there are three problems that often arise with this approach: 1) Random sampling with rejection may be unlikely to yield experiment sequences that satisfy the set of desired constraints. For example, this issue is particularly problematic in designs

in which researchers attempt to account for sequential (across-trial) dependencies between experiment factors. 2) Even if a solution can be found, simple forms of randomization do not ensure that the samples are unbiased. This is of particular concern when the factors are not independent of one another, and especially when there are sequential dependencies. This can greatly complicate, and in some cases seriously constrain how subsets of factors can be crossed with one another, introducing biases in the sampling process that can lead to unintended (and often unrecognized) confounds in the design. 3) Even when randomization is executed in a way that ensures an unbiased design, it is often done using custom-written code that can be difficult to understand and therefore replicate, and is almost always specific to the particular experiment and thus not more generally useful.

As an example, consider a Stroop experiment (Stroop, 1935) in which participants are asked to name the color in which a color word is displayed (e.g., say “green” to the written word RED displayed in green). A simple design may seek to counterbalance just the combination of words and colors to be used (i.e., by insuring that every word appears in every color). However, a more complicated design may demand that the color and word disagree (e.g., the word RED displayed in green) for every trial; that the color and word used on a given trial be different than the ones used on the trial before and after it; or that the task to be performed (i.e., reading the word vs. naming the color) be switched a certain number of times. Such requirements have become increasingly common in cognitive psychological studies, as researchers seek to implement more sophisticated designs that more precisely address processes of interest (e.g., the mechanisms responsible for switching attention between tasks). Constructing algorithms that satisfy the constraints of such designs, while uniformly sampling (i.e., counterbalancing) all legal variants of the design, can quickly become a challenge even for a professional programmer. However, failure to do so may either not yield the desired design and/or introduce unintended biases. Indeed, as designs have become more complex, so too has the frequency of confounds identified in such designs (e.g. Allport & Wylie, 1999; Cooper & Marí-Beffa, 2008; Jou, 2014; Logan & Schneider,

2010; Zmigrod & Hommel, 2013). Furthermore, even when appropriate algorithms are used, the code that implements them is usually complex and design-specific, posing problems for interpretability and reproducibility (Miłkowski, Hensel, & Hohol, 2018; Peng, 2011; Sochat et al., 2016).

A New Language for Experimental Design

To address these concerns, we have developed the first generation of a software tool that is designed to narrow the gap between the specification of an experimental design and its implementation in trial-generating programs. Our approach is to create a declarative language implemented in **Python** for describing experimental designs in factorial form, and then express the design in a format that can be used by standard computational sampling algorithms to ensure—as best as possible—that the trial sequences generated are as unbiased as possible. That is, instead of separately describing the intended design and then writing down a set of **for** loops and calls to **random** that are intended to implement that design, an experimenter can describe the design precisely in a language that is tailored to the problem of experimental design. The generation of trials from that description can then be fully automated, insuring that it is done in a consistent and well-defined way, and with the hope of providing increasingly strong guarantees about the statistical properties of generated trials as the tool is improved by incorporating new advances in sampling and sample analysis algorithms, and/or computational power increases.

The language we describe, **SweetPea**, builds on the foundation of factor-based design. In addition to independent factors that can be crossed in various ways, the language adds a notion of *derived factors*, which allow an experimenter to give names to properties of interest (e.g., whether two properties of a stimulus are “congruent” or “incongruent”). Finally, the experimenter can impose explicit *constraints* on the factors, derived factors, and even the relationship of different factors within a window of adjacent trials in the experiment. This combination of *factors* and *constraints* directly addresses the growing complexity of experimental designs in empirical sciences and

machine learning. Generating a run of the experiment corresponds to sampling from the solution space of those constraints, ideally with a guarantee of uniformity that avoids accidental correlations in generated trials.

The implementation of SweetPea mirrors the design goal of separating the “what” of experimental design from the “how” of generating trials. Our language front-end converts an experimental design into a constraint-solving problem, where various constructs imply constraints on trial content and order. The constraint problem then can be delegated to a combination of custom and existing solvers and samplers that act as the back-end of the implementation. This structure substantially decouples the problem of creating an expressive experiment-design language from the problem of implementing it efficiently enough, so researchers can explore both directions concurrently.

The SweetPea language is an open-source project² and, as such, work in progress. The development of SweetPea focused on two fronts. First, with respect to expressiveness, while the language has proven rich enough for several designs of immediate interest (e.g., for which generating trials by one-off programs has been challenging; see Section *Walkthrough Examples*), it is still limited in a number of ways. For example, it currently does not support the weighted sampling of factor levels and weighted crossings thereof. Our hope is that continued development (on our part, and through community support), will refine and expand the language and its semantics to cover a greater range of experimental designs. Second, with respect to sampling guarantees, we have not yet found a practical solution to the problem of guaranteeing a uniform sample of the space of all possible constraint solutions (valid trial sequences); that general problem is an area of active research in computer science, and we report on our experience to date using currently available solutions. However, by standardizing the expression of experimental designs, and its interface with sampling methods, our tool ensures that as advances are made in computational algorithms for sampling and sample analysis, these can be easily introduced into SweetPea, thus allowing

² We invite contributions to SweetPea’s open-source repository (<https://github.com/sweetpea-org/>).

experimental design methods to readily leverage these advances.

Using SweetPea

SweetPea is a “language” in the sense that it defines a vocabulary and set of abstractions for describing experiments, but it is implemented as a library in Python. An experiment description in SweetPea is a Python value that is constructed with lists, objects, and primitive data types (e.g. strings or integers) that represent factor names and level values. Passing an experiment description to SweetPea’s `synthesize_trials` function generates a list of runs. Each run in that result is reported as a dictionary mapping factor names to level values, with one level value for each trial in the run. The interpretation of factor and level strings (e.g., interpreting a “color” factor’s “red” level to mean that a stimulus should be displayed in the color red) is outside the scope of SweetPea, but embedding SweetPea in Python simplifies the bridge between trial generation and presentation, since experiment-presentation systems, such as PsychoPy (Peirce, 2007, 2009) or Expyriment (Krause & Lindemann, 2014), are also available as Python libraries. Similarly, SweetPea can interface with data processing pipelines used to train artificial systems in Python, such as PyTorch (Paszke et al., 2019) and TensorFlow (Abadi et al., 2016), as well as simulation environments targeted specifically at cognitive and neuroscientific models such as PsyNeuLink³.

Defining Factors

Figure 1 shows the complete grammar of SweetPea experiment descriptions in terms of Python values. For example, the production `[block, ...]` means a list of *blocks* values, and the production `factor(factor_name, [level_value, ...])` means an object produced by passing a *factor_name* and a list of *level_values* to the `factor` function.

³ <http://www.PsyNeuLink.org>

<i>experiment</i>	=	<i>block</i> [<i>block</i> , ...]
<i>block</i>	=	<code>fully_cross_block(design, crossing, constraints)</code>
<i>design</i>	=	[<i>factor</i> , ...]
<i>crossing</i>	=	[<i>factor</i> , ...]
<i>constraints</i>	=	[<i>constraint</i> , ...]
<i>factor</i>	=	<code>factor(factor_name, [level_value, ...])</code> <code>factor(factor_name, [derived_level, ...])</code>
<i>derived_level</i>	=	<code>derived_level(level_value, derivation)</code>
<i>derivation</i>	=	<code>within_trial(predicate, [factor, ...])</code> <code>transition(predicate, [factor, ...])</code> <code>window(stride, width, function, [factor, ...])</code>
<i>predicate</i>	=	a function from [<i>level_value</i> , ...] to <i>boolean</i>
<i>constraint</i>	=	<code>exclude(factor_name, level_value)</code> <code>at_most_k_in_a_row(k, levels)</code> <code>exactly_k_in_a_row(k, levels)</code> <code>no_more_than_k_in_a_row(k, levels)</code> <code>minimum_trials(trial_number)</code>
<i>levels</i>	=	(<i>factor_name</i> , <i>level_value</i>) <i>factor</i>
<i>factor_name</i>	=	a string
<i>level_value</i>	=	a primitive data type (string, integer, float, boolean)
<i>k, stride, width</i>	=	a positive integer

Figure 1. SweetPea grammar.

An experiment consists of one or more blocks, where each block is a sequence of trials, and each trial combines one level from each of the block's factors. A factor is either *independent* or *derived*, depending on whether its levels are primitive data types (e.g. string or integer) or created by the `derived_level` function. For example,

consider the Stroop experiment mentioned above, in which participants are asked to name the color in which a color word is displayed. Each trial of the Stroop experiment consists of the independent factors `color` and `word` that can be defined as

```
color_list = ["red", "green", "blue", "brown"]
color = factor("color", color_list)
word = factor("word", color_list)
```

Note that `color` here is a Python variable, while `"color"` is a factor name. A SweetPea experiment description has factor names, not variables, but Python variables are useful to give names to parts of SweetPea experiments.

In a simple variant of the Stroop experiment, it may be sufficient to fully cross the independent factors `color` and `word`, and then randomize the order in which the trials are drawn from the crossing to produce a block of trials. In other cases, the experimenter may want to constrain the occurrences of “congruent” trials (i.e., in which the color and word match) versus “incongruent” trials (in which they are different). This notion of congruence can be made explicit by deriving it in terms of the `color` and `word` factors:

```
congruent = derived_level("con", within_trial(operator.eq, [color, word]))
incongruent = derived_level("inc", within_trial(operator.ne, [color, word]))
congruence = factor("congruence", [congruent, incongruent])
```

A derived level such as `congruent` or `incongruent` applies a predicate to levels drawn from other factors. The levels might be candidates within a trial or across trials; the `within_trial` function applies the predicate to candidate levels within a single trial, which is what is needed for defining congruence in the example above. The predicate given to `within_trial` must take the same number of arguments as factors in a list given to `within_trial`. In the example, the Python predicate `operator.eq` expects two arguments and returns `True` when they are the same, while `operator.ne` returns `True` when two arguments are different, so those predicates make sense given levels from `color` and `word` as specified by `[color, word]`. Thus, the uses of `within_trial` for `congruent` and `incongruent` select complementary portions of the color–word crossing. The `derived_level` function assigns a name to each of those

spaces, and the resulting levels are combined with **factor** to define a **congruence** factor. The **factor** function requires that each derived level identifies a distinct portion of a crossing and that every part of the crossing is covered by one of the levels.

If **color**, **word**, and **congruence** are declared to be the factors of a block, then each trial in that block will have a level for each of those three factors. For example, two possible trials are

```
{ 'color': 'red', 'word': 'green', 'congruence': 'inc' }
{ 'color': 'red', 'word': 'red', 'congruence': 'con' }
```

SweetPea's `print_experiments` function can be used to print these Python dictionaries more compactly:

```
color red | word blue | congruence inc
color red | word red  | congruence con
```

The following is *not* a possible trial, because the derived factor's level is not consistent with the other factors' levels:

```
color red | word blue | congruence con
```

While the **congruence** level could be inferred from the **color** and **word** levels, a benefit of explicitly defining **congruence** is that a block can then be specified to be the crossing of, say, **color** and **congruence**, which ensures that the block contains every color once with each congruence. A possible block produced by `synthesize_trials` in that case is

```
color red   | word green | congruence inc
color red   | word red   | congruence con
color blue  | word blue  | congruence con
color green | word green | congruence con
color green | word red   | congruence inc
color blue  | word red   | congruence inc
```

Note that this sequence does not include every possible **color-word** combination, as would be produced by instead crossing **color** and **word**. Crossing **word** and **congruence** would also work, and would ensure that each word appears twice within the block. Crossing all three factors would *fail* and cause `synthesize_trials` to report an error,

because some combinations of levels violate the constraint inherent in the definition of **congruence**, such as

```
color red   | word blue   | congruence con
```

Crossing and Constraints

In addition to the constraints that are inherent to a derived factor, an experimental design can also impose other constraints on generated trials. Those may rule out particular combinations of levels (i.e., portions of a crossing), or they may constrain the order of combinations. Typically and most usefully, some aspects of the trials and ordering can remain unconstrained. For example, in the cases above the order of trials within a crossing is unconstrained, so that **synthesize_trials** can randomly select among all possible orderings. Similarly, crossing **color** and **congruence** allows the **word** level to vary randomly among the two possibilities for each incongruent trial.

The general constructor to describe a block of trials is **fully_cross_block**, which expects three arguments:

- *design*: A list of all factors in the experiment. Every trial in the block will have a single level value for each of these factors.
- *crossing*: A subset of the factors in *design* that are to be fully crossed in each run of the block. This crossing imposes two constraints on the result: the minimum number of trials per run is the product of the factor sizes, and each of those runs must have the same number of instances for every distinct combination of levels among the factors.
- *constraints*: Additional constraints on either (a) the minimum number of trials, (b) the order of the *crossing* combinations, or (c) the level of factors in *design* that are not in *crossing*.

As an example constraint, suppose that we want to fully cross **color** and **word** in our Stroop experiment, but we do not want two consecutive congruent trials. The **at_most_k_in_a_row** function lets us express this constraint:

```
one_con_at_a_time = at_most_k_in_a_row(1, (congruence, congruent))
```

The `at_most_k_in_a_row`, `exactly_k_in_a_row`, and `no_more_than_k_in_a_row` functions take a count for “k” and one or more levels. A single level is specified as a tuple of a `factor_name` and `level_value`, while a factor argument indicates all of the levels of the factor, effectively duplicating the constraint for each level.

Assembling the factors and constraints into a block completes the design of a single-block experiment:

```
design          = [color, word, congruence]
crossing        = [color, word]
constraints     = [one_con_at_a_time]

experiment     = fully_cross_block(design, crossing, constraints)
```

A potential run of this experiment is

```
color blue   | word green | congruence inc
color green  | word blue  | congruence inc
color blue   | word blue  | congruence con
color blue   | word red   | congruence inc
color green  | word red   | congruence inc
color red    | word red   | congruence con
color red    | word green | congruence inc
color green  | word green | congruence con
color red    | word blue  | congruence inc
```

Note that trying to constrain the experiment to at most one incongruent trial in a row would fail, because there are more incongruent pairs than congruent pairs in a color–word crossing, so adjacent incongruent trials are inevitable. SweetPea’s `synthesize_trials` function reports an error when constraints have no solution, but it cannot always provide a simple explanation for why an experiment is over-constrained.

Transitions and Windows

The `one_con_at_a_time` derived factor sets up a kind of transition constraint, in that a congruent trial must transition to an incongruent trial. SweetPea also supports a

more direct and general way of controlling transitions. For example, suppose that each trial of our experiment should be related to the previous trial by having either the same color or the same text, but not both. By using `transition` instead of `within_trial`, we can create a derived factor that categorizes a trial as either changing exactly one independent factor or both, relative to the previous trial:

```
def one_diff(colors, words):
    if (colors[0] == colors[1]):
        return words[0] != words[1]
    else:
        return words[0] == words[1]

def both_diff(colors, words):
    return not one_diff(colors, words)

one = derived_level("one", transition(one_diff, [color, word]))
both = derived_level("both", transition(both_diff, [color, word]))
changed = factor("changed", [one, both])
```

Like `within_trial`, the predicate provided to `transition` must accept one argument for each factor listed to `transition`. But while `within_trial` gets a level value for each argument, `transition` gets a list of two level values for each argument. The first value in the list is the previous trial's value, and the second element in the list is the current trial's value. The `one_diff` predicate here determines whether exactly one of the lists has the same values for its two elements.

If we just add `changed` to our experiment, then the result shows whether one of color and word change or both:

```
color blue   | word green | congruence inc
color green  | word blue  | congruence inc | changed both
color blue   | word blue  | congruence con | changed one
color blue   | word red   | congruence inc | changed one
color green  | word red   | congruence inc | changed one
color red    | word red   | congruence con | changed one
color red    | word green | congruence inc | changed one
```

```

color green | word green | congruence con | changed one
color red   | word blue  | congruence inc | changed both

```

Notice that there is no `changed` value for the first trial, since a trial needs a preceding trial to determine its transition type. To rule out trials that change both `color` and `word` relative to the previous one, we can add an `exclude` constraint:

```

constraints = [one_con_at_a_time,
               exclude(changed, both)]

```

As it happens, the `one_con_at_a_time` constraint is now redundant, since a congruent trial must be followed by one that changes one of `color` or `word`. Such redundancies are not always obvious, however, and including them in the design causes no problems. With this new constraint, `synthesize_trials` cannot generate the previous run, and it might instead generate this one:

```

color green | text red   | congruence inc
color green | text green | congruence con | changed one
color red   | text blue  | congruence inc | changed one
color red   | text red   | congruence con | changed one
color blue  | text red   | congruence inc | changed one
color blue  | text blue  | congruence con | changed one
color green | text blue  | congruence inc | changed one
color blue  | text green | congruence inc | changed one
color red   | text green | congruence inc | changed one

```

Derived factors defined by `transition` can also be included in an experiment's crossing. For example, if we define the experiment's crossing to be `color` and `changed` and omit further constraints, then a possible experiment is

```

color red   | text red   | congruence con | changed
color green | text green | congruence con | changed both
color blue  | text green | congruence inc | changed one
color blue  | text green | congruence inc | changed both
color red   | text blue  | congruence inc | changed both
color green | text blue  | congruence inc | changed one
color red   | text blue  | congruence inc | changed one

```


In this experiment, each `color` is included once for each `changed`, which is why the experiment includes six trials to cover all combinations, plus one more at the start to enable the `changed` transition. Not every `color` and `word` combination appears, since `word` has been removed from the crossing. If we cross all three of `color`, `word`, and `changed`, the resulting experiment would have $3 \times 3 \times 2 + 1 = 19$ trials.

The `within_trial` and `transition` functions for derived levels are both shorthands for a general `window` constructor. The general `window` form supports relationships among any fixed number of trials, and it allows the trials to be separated by a given stride within the run instead of requiring them to be adjacent. We illustrate the use of the `window` function in the walkthrough example *Designing a 2-Back Task*.

Solving Experimental Designs

Our simple Stroop variant has only 9 trials in a full crossing of `color` and `word`. Those 9 trials can be ordered in $(3 \times 3)! = 362,880$ ways. For that scale, a practical implementation could simply enumerate all permutations, filter ones that do not conform to a constraint, and then randomly select from the rest. Raising the number of color names from 3 to 5, however, creates 25 elements in the crossing, which means 1.5×10^{25} permutations; it's still a tiny experiment, but generate-all-and-filter is not remotely practical. For sufficiently constrained experiments, even generate-one-and-test is not practical, because the generated experiment is unlikely to satisfy the constraints.

The SweetPea implementation can avoid generate-and-test and instead use an external program to find a trial sequence that crosses all specified factors and that satisfies all indicated constraints (Cherkaev, 2019). The external program is a *SAT solver*, which can find a solution to a general logical formula (or report that no such solution exists). In principle, finding a solution to an arbitrary logical formula is intractably difficult. In practice, modern SAT solvers have heuristics that efficiently find solutions. SweetPea compiles the constraints that define an experiment into a disjunctive normal form—a logical expression that a SAT solver can recognize—and then it translates the SAT solver's answer back to experiment terms.

Using a SAT solver is good for finding *one* valid run of an experiment. A second run can be generated by further constraining the solution so that it’s different from the first solution. While this strategy can be used to generate multiple experiments, it unfortunately fails to solve one problem raised above, which is to sample the space of solutions uniformly; new solutions generated by simply ruling out previous solutions might explore only a small part of the solution space.

A *SAT sampler* can drive a SAT solver to explore the full space of solutions. SweetPea uses Unigen2 (Chakraborty, Meel, & Vardi, 2014) to implement uniform sampling of small experimental designs.⁴ Our attempt to use SAT samplers for medium-to-large designs has not so far been effective, however. The problem is that current SAT samplers rely on a formula having a certain structure, and the way that SweetPea currently encodes experiments does not fit that structure. Indeed, the mismatch appears to be fundamentally related to the problem of generating permutations of possible trials, which suggests that no encoding will work.

SAT solving and sampling can work for some experimental designs, generate-and-test can work for some designs, and strategies that involve sampling from numbers with a bijection between numbers and solutions can work in other designs. By framing the problem in terms of a declarative specification of an experiment, SweetPea separates the problem of describing an experiment and generating the experiment—much the same way that SAT solvers separate the problem of specifying formulas and solving those formulas.

Much of our ongoing work is about finding new implementation strategies to “solve” experiments and to automatically determine which strategies will work for a given experiment description. For now, SweetPea offers a sampling strategy that can handle experiments without constraints by just picking a random permutation among the crossing. SweetPea also currently offers a strategy for finding trial sequences for small-to-medium designs, which can be found using a SAT solver without a sampler, but it does not provide a guarantee of uniform sampling. New samplers will likely

⁴ In future releases of SweetPea, we will leverage Unigen4 (<https://github.com/meelgroup/unigen>).

involve a hybrid of these approaches.

Walkthrough Examples

This section illustrates the application of SweetPea to three different experimental design problems. The first design problem (Stroop task) illustrates the use of basic and derived factors, the second design problem (task switching) introduces factors that describe transitions between two trials, and the third design problem (2-Back task) requires solving sequential dependencies between more than two trials. We chose to focus on experiments that are commonly used to study cognitive control (i.e., the ability to flexibly pursue goal-directed behavior in the face of distraction). However these experiments share similarities with experiments in other domains of psychology, as well as other empirical disciplines, such as neuroscience and machine learning. Before walking through each example, we describe the installation of SweetPea. We then focus on each experiment, by first describing the experimental design problem, then showing the code for expressing and solving the problem in SweetPea, and finally walking the reader line-by-line through the example code. In addition, each example includes a solution for the design problem that was identified by SweetPea. Due to their medium complexity, the experiments described in this section are not amenable to uniform sampling, so we sampled experiment solutions non-uniformly from the space of possible solutions. The reader may find examples that implement uniform sampling, as well as other examples—both proof of concept and for realistic experimental designs—online (<https://sweetpea-org.github.io/>).

Installation of SweetPea

The current version of SweetPea requires Python 3.6 or later. It relies on a backend server to solve experimental designs. The backend server requires docker (www.docker.com; Merkel, 2014) to be installed and running. SweetPea itself can be installed in a local Python environment using

```
pip install sweetpea
```

Designing a Stroop Task

In previous sections we described the Stroop task (Stroop, 1935) in which participants are typically asked to name the color in which a color word is displayed (e.g., say “green” to the word **RED** displayed in green). Here, we consider the Stroop color naming experiment with the two regular factors described above: the color factor representing the color in which the stimulus is displayed, and having four levels: red, green, blue, brown; and the word factor representing the word itself, also having four levels corresponding to each of the colors (Fig. 2A). On a given trial, one level from each factor is used to generate the stimulus, and the participant is required to respond to the color factor, indicating the color in which the word was displayed. For instance, they may be required to press the left arrow key if the word was displayed in red, the right arrow key if it was in green, the up arrow key if it was in blue, the down arrow if it was in brown. This results in a response factor with four levels: left, right, up and down.

Figure 2. Experimental designs used in examples. Each experiment can be described as a sequence of trials. In each of the three examples, a trial consists of one or two displays of a stimulus that appear in sequence. The figure shows experiment factors (in bold) with their respective levels (regular font) describing each trial. The reader may refer to the text for a detailed description of each experiment. (A) Stroop color naming experiment (Stroop, 1935). Participants are asked to indicate the color in which a color word is displayed, in this case with a button press. Relevant experiment factors include the color of the word, the word itself, the correct response button associated with the color, as well as stimulus congruency. A stimulus is considered congruent if the color and word are the same, and incongruent otherwise. (B) Cued task switching experiment (Meiran, 1996; Sudevan & Taylor, 1987). The same stimuli are used as in the Stroop experiment, however in this case participants are instructed to press a button corresponding either to the color of the stimulus (color naming) or to the word (word reading). Each trial begins with the display of a shape cue that signals which task to perform. A circle indicates that the color naming task should be performed for the ensuing stimulus, and a square indicates that the word reading task should be performed on that trial. Each trial (except the first) can be characterized as a transition between tasks relative to the previous trial: Participants may either repeat the same task or switch to a different task. The same applies to transitions between responses from one trial to the next. (C) 2-Back task (Cohen et al., 1997). Participants are presented with a sequence of letters, and instructed to press a button if the letter on the current trial matches the letter presented two trials back. The letter on the current trial is designated as a “target” if it matches in this way, and is considered to be a “lure” if it (a) is not a target and (b) matches the letter on the previous trial.

In the general case, an experimenter may want to ensure that each color is paired with each word. This could be achieved by crossing all colors with all words. The full crossing of colors and words includes conditions in which the color and the word are the same (congruent; e.g. the word **RED** displayed in red), as well as conditions in which they are different (incongruent; e.g. the written word **GREEN** displayed in red). Together, the two conditions define the congruency of a trial. In this particular experiment, we wish to pair each color with each word, subject to the constraint that only incongruent trials are included. The code shown in Listing 1 generates a minimum of 20 trials of this experiment.

```

1 from sweetpea.primitives import factor, derived_level, within_trial
2 from sweetpea.constraints import minimum_trials, exclude
3 from sweetpea import fully_cross_block, synthesize_trials_non_uniform, print_experiments
  , experiment_to_csv
4
5 # color and word factors
6
7 color      = factor("color",  ["red", "green", "blue", "brown"])
8 word       = factor("word",  ["red", "green", "blue", "brown"])
9
10 # response factor
11
12 def is_response_left(color):
13     return color == "red"
14 def is_response_right(color):
15     return color == "green"
16 def is_response_up(color):
17     return color == "blue"
18 def is_response_down(color):
19     return color == "brown"
20
21 response = factor("response", [
22     derived_level("left", within_trial(is_response_left, [color])),
23     derived_level("right", within_trial(is_response_right, [color])),
24     derived_level("up", within_trial(is_response_up, [color])),
25     derived_level("down", within_trial(is_response_down, [color]))
26 ])
27
28 # congruency factor
29
30 def is_congruent(color, word):

```

```

31     return color == word
32
33 def is_incongruent(color, word):
34     return not is_congruent(color, word)
35
36 congruent = derived_level("congruent", within_trial(is_congruent, [color, word]))
37 incongruent = derived_level("incongruent", within_trial(is_incongruent, [color, word]))
38
39 congruency = factor("congruency", [
40     congruent,
41     incongruent
42 ])
43
44 # constraints
45
46 trial_constraint = minimum_trials(20)
47 exclusion_constraint = exclude(congruency, congruent)
48
49 constraints = [trial_constraint, exclusion_constraint]
50
51 # experiment
52
53 design      = [color, word, response, congruency]
54 crossing    = [color, word]
55 block      = fully_cross_block(design, crossing, constraints,
56                               require_complete_crossing=False)
57
58 experiments = synthesize_trials_non_uniform(block, 1)
59
60 print_experiments(block, experiments)
61
62 experiment_to_csv(experiments, file_prefix="experiment")

```

Listing 1: Example of a Stroop task.

Lines 1-3 import the SweetPea modules. Modules in line 1 are needed to specify regular and derived factors. Modules in line 2 are required to implement design constraints, such as specification of the minimum number of trials and the exclusion of factor levels. The imported modules in line 3 serve to generate, sample, print and export the final experiment.

Lines 7-8 define the two basic factors `color` and `word`. Levels of the `response` factor are dependent on the factor `color`. Each of its derived levels is defined by a

predicate that takes as input the `color` factor describing the current stimulus (lines 12-19). For instance, the following function is used to define a leftward response:

```
def is_response_left(color):
    return color == "red"
```

The function returns true if the factor `color` on a given trial evaluates to the level `"red"`. The function is used to specify the derived level `"left"` in the definition of the factor `response` (lines 21-26):

```
derived_level("left", within_trial(is_response_left, [color])),
```

The first argument specifies the level name `"left"`. The second argument is the function `within_trial` described in Section *Defining Factors*. It passes the `color` factor of the current trial to the `is_response_left` predicate, to determine whether the factor `response` evaluates to the level `"left"`. The `congruency` factor is defined in a similar fashion (lines 30-42). Both levels of this factor are dependent on the `color` and `word` factors of the current trial. Thus, the predicates specifying these levels take the `color` and the `word` factors as arguments (lines 30-34). For instance, a trial is considered congruent if the color and word of the current trial match:

```
def is_congruent(color, word):
    return color == word
```

Note that the definition of the corresponding derived level must include `within_trial`, with both `color` and `word` as arguments (line 36):

```
derived_level("congruent", within_trial(is_congruent, [color, word]))
```

The experiment code specifies two constraints. First, it specifies the minimum number of trials (line 46). The `minimum_trials` constraint ensures that the experiment sequence includes at least 20 trials. Note that a full crossing of all valid trials (without 4 possible congruent trials) requires a multiple of $4 \times 4 - 4 = 12$ trials. Thus, `minimum_trials` will not satisfy a full crossing unless the specified minimum number of trials is a multiple of 12. To mitigate this issue, SweetPea successively samples trials without replacement from counterbalanced blocks. In this example, the first 12 trials are sampled from a block of 12 counterbalanced trials, and the remaining 8 trials are

sampled without replacement from another counterbalanced block.⁵ The second constraint defines an exclusion criterion according to which the level **congruent** of the factor **congruency** is excluded from the experiment (line 47). All factors to be included in the design are listed in line 53. The crossing between all colors and all words is specified in line 54. The design, crossing and constraints are used to define a fully crossed experiment block (subject to said constraints; lines 55-56). However, a complete crossing between all colors and words is not possible because we want to exclude all congruent trials for which the color and the word match. Thus, **fully_cross_block** would return no solution to the experimental design unless we allow the crossing to be incomplete, by setting **require_complete_crossing** to **False**.

Line 58 specifies how the experiment should be generated:

```
experiments = synthesize_trials_non_uniform(block, 1)
```

In this case, we sample the experiment block only once, thus the argument 1. The function **synthesize_trials_non_uniform** solves for experiment sequences without guaranteeing that they are sampled uniformly from the space of all possible solutions (see Section *Solving Experimental Designs*). Finally, we print the experiment in line 60, and export it to a comma-separated values (CSV) file named “experiment_0.csv”⁶ in line 62. An example output is shown in Table 1.

Designing a Task Switching Experiment

In many experiments, the sequence in which trials are presented is an important part of the design. For example, in task switching paradigms—used to study the flexibility with which people can adapt their behavior—the transition between trial

⁵ The current version of SweetPea only supports this counterbalancing scheme for experimental designs without transition factors and constraints. The counterbalancing described in the text is done independently for each experimental design. Future versions of SweetPea will coordinate this counterbalancing scheme across multiple experiment sequences, e.g. to allow for the counterbalancing of remaining trials across participants (Latin square design).

⁶ Both the **print_experiments** and the **experiment_to_csv** functions require a list of experiments. Here, we want to generate just one experiment, so both functions will output one experiment sequence.

types is an important factor. One example of this is the cued task switching paradigm, in which participants receive a task cue on every trial that instructs them which of two (or more) tasks they should perform on the current trial (Fig. 2B). The cue may instruct them either to repeat the task from the previous trial (task repetition) or to switch to a different task (task switch). A common measure in such designs is the cost associated with task switches—that is, reaction time or accuracy on switch relative to repetition trials. Thus, the design must consider task transition as a factor.

Here, we consider a task switching paradigm that builds on the Stroop experiment described in the previous example. In this experiment, participants are presented with a sequence of Stroop stimuli (i.e., color words displayed in a particular color). For simplicity, we consider only two colors (red and green), resulting in two levels for the color factor and two levels for the word factor. The task factor indicates which of the two tasks the participant is instructed to perform on a given trial, with two levels: color naming (respond to the color in which the word is displayed) and word reading (respond to the word). Thus, the correct response on every trial depends on both the stimulus and the relevant task. As in the previous example, we assume that the participant responds by pressing one of two buttons, but the stimulus-response mapping is the opposite for the two tasks: for color naming, the left button should be pressed for the color red and the right button for the color green; conversely, for word reading, the right button should be pressed for the word “red” and the left button for the word “green”. Finally, there is a task transition factor, with two levels: repeat (if the instructed task for the current trial is the same as the last trial) and switch (if it is different).

In addition to the factors described above, it may also be important to include a response transition factor, that determines whether the response required on the current trial (i.e., whether the left or right button is the correct response) is the same or different as the one required on the previous trial. For example, this response transition factor has been shown to interact with performance costs associated with task switches (Kiesel et al., 2010; Rogers & Monsell, 1995). Thus, to control for this, it may be important to ensure that the same number of response transitions occurs in the task

switch and task repetition conditions. This can be done by specifying a full crossing of the task transition, response transition, color, word and task factors.

Fully crossing the factors described above will ensure a balanced sampling of all combinations of their levels, including types of task and response transitions, over the course of the experiment. However, when selecting trials to execute, it does not necessarily preclude the possibility of undesired local sequences, or “runs”—that is, sequences of trials in which some level of a factor (e.g., whether the current trial is a task repetition or response repetition) remains the same, or alternates in some seemingly predictable but undesired way—for several trials in a row. This could invite misleading expectations from the participant, or otherwise impact performance in undesired ways. While these could also be specified as higher level factors, it can be more convenient to control for this by imposing constraints on the number times a particular level of a given factor, or a particular crossing of two or more factors is allowed to occur in sequence. As an example, here we limit task transitions to four of the same type, and the number of consecutive response repetitions or response switches to four. The code shown in Listing 2 implements this design.

```

1 from sweetpea.primitives import factor, derived_level, within_trial, transition
2 from sweetpea.constraints import no_more_than_k_in_a_row
3 from sweetpea import fully_cross_block, synthesize_trials_non_uniform, print_experiments
4
5 # color, word and task factors
6
7 color      = factor("color", ["red", "green"])
8 word       = factor("word", ["red", "green"])
9 task       = factor("task", ["color naming", "word reading"])
10
11 # congruency factor
12
13 def is_congruent(color, word):
14     return color == word
15
16 def is_incongruent(color, word):
17     return not is_congruent(color, word)
18
19 congruency = factor("congruency", [
20     derived_level("congruent",

```

```

21         within_trial(is_congruent, [color, word])),
22     derived_level("incongruent",
23         within_trial(is_incongruent, [color, word]))
24 ])
25
26 # response factor
27
28 def is_response_left(color, word, task):
29     return (task == "color naming" and color == "red") or \
30         (task == "word reading" and word == "green")
31 def is_response_right(color, word, task):
32     return (task == "color naming" and color == "green") or \
33         (task == "word reading" and word == "red")
34
35 response = factor("response", [
36     derived_level("left",
37         within_trial(is_response_left, [color, word, task])),
38     derived_level("right",
39         within_trial(is_response_right, [color, word, task]))
40 ])
41
42 # task transition factor
43
44 def is_task_repetition(task):
45     return task[0] == task[1]
46
47 def is_task_switch(task):
48     return not is_task_repetition(task)
49
50 task_transition = factor("task_transition", [
51     derived_level("repetition",
52         transition(is_task_repetition, [task])),
53     derived_level("switch",
54         transition(is_task_switch, [task]))
55 ])
56
57 # response transition factor
58
59 def is_response_repeat(response):
60     return response[0] == response[1]
61
62 def is_response_switch(response):
63     return not is_response_repeat(response)
64

```

```

65 response_transition = factor("response_transition", [
66     derived_level("repetition",
67         transition(is_response_repeat, [response])),
68     derived_level("switch",
69         transition(is_response_switch, [response]))
70 ])
71
72 # constraints
73
74 constraints = [
75     no_more_than_k_in_a_row(4, task_transition),
76     no_more_than_k_in_a_row(4, response_transition)]
77
78 # experiment
79
80 design      = [color, word, task, congruency, response,
81                task_transition, response_transition]
82 crossing    = [color, word, task, task_transition, response_transition]
83 block       = fully_cross_block(design, crossing, constraints)
84
85 experiments = synthesize_trials_non_uniform(block, 1)
86
87 print_experiments(block, experiments)

```

Listing 2: Example of a task switching experiment.

Lines 1-3 import the SweetPea modules needed to specify the experiment as described for the previous example. We also include `transition` (line 1) to define factor levels based on transitions between trials.

Lines 7-9 define the three regular factors, `color`, `word` and `task`. Lines 13-24 define the derived factor `congruency`, as in the previous example (cf. Listing 1, lines 30-42). However, for conciseness, here the `derived_level` function is directly passed to `factor`.

Lines 28-40 define the derived factor `response`. The levels of this factor are defined by the two functions `is_response_left` and `is_response_right`. Note that each of the functions depends on the three factors `color`, `word` and `task`. For instance, the `is_response_left` function

```

def is_response_left(color, word, task):
    return (task == "color naming" and color == "red") or \

```

```
(task == "word reading" and word == "green")
```

implements the rule that the left response button should be pressed if the task is "color naming" and the color is "red", or if the task is "word reading" and the word is "green".

The `task transition` factor is defined in lines 44-55. The levels of this factor are dependent on the factor `task` in the previous trial and the current trial. The predicate `is_task_repetition` expresses this between-trial dependency to define the task repetition level:

```
def is_task_repetition(task):
    return task[0] == task[1]
```

In this example, the factor `task` is passed as an ordered list with two elements. The first and second elements of the list encode the task of the previous trial (`task[0]`) and the current trial (`task[1]`), respectively. If the two are the same, the current trial is considered a task repetition, and a task switch otherwise. Note that the derived levels of `task_transition` are now defined using the `transition` function. For instance,

```
derived_level("repetition",
              transition(is_task_repetition, [task])),
```

defines the level "repetition", and uses the `transition` function to pass the factor `task` as a list (encoding the task on the current and previous trial) to the predicate `is_task_repetition`. Lines 59-70 define `response_transition` in an analogous manner.

Lines 74-76 define constraints for the experiment. The function `no_more_than_k_in_a_row` in line 75 implements the sequential constraint that each level of the factor `task_transition` cannot occur more than four times in a row; the same is declared for the factor `response_transition` in line 76.

Analogous to Listing 1, the design, the crossing, as well as the constraints are integrated into an experiment (lines 80-83) that is first sampled non-uniformly (line 85), and then printed (line 87).

The output shown in Table 2 illustrates one solution to the specified design. Note

that the factors `task_transition` and `response_transition` are not defined for the first trial of the experiment, simply because there exists no preceding trial with the factors `task` and `response`, respectively. Full counterbalancing of five factors, each of which has two levels, requires at least $2^5 = 32$ trials. SweetPea adds an additional filler trial at the beginning of the experiment sequence to accommodate the circumstance that transition factors cannot be defined for the first trial of an experiment.

Designing a 2-Back Task

Some experiments may involve factors that are dependent on more than two consecutive trials. The function `window` allows the user to define such factors in SweetPea. Here, we illustrate its functionality in the design of an N-Back task—a psychological task commonly used to assess working memory performance, e.g. how well participants can update and maintain task-relevant information over time. In the N-Back task considered here, participants are presented with a sequence of letters, one letter per trial (Cohen et al., 1997). Participants are instructed to press a button if the letter on the current trial matches the letter some number N trials back. For simplicity, we consider a 2-Back task in which participants should press a button if the current letter was presented $N = 2$ trials back (Fig. 2C). The letter factor describes which of the following letters is presented on the current trial: “A”, “B”, “C”, “D”, “E”, “F”. The target factor determines whether the letter on the current trial matches the letter two trials back, in which case the participant has to press the button. Finally, we consider a trial to be a lure if (a) it isn’t a target trial but (b) the letter on the current trial matches the letter one trial back. Lure trials are important because they can help determine whether, when participants make mistakes, it is because they are having trouble remembering the letters themselves (e.g., a problem with maintenance) or the order in which they were presented (e.g., a problem with updating). It is also possible that participants may be better at responding to some letters than others. Measures of working memory performance could be biased if those letters occur more often in the experiment. This bias can be avoided by ensuring that each letter is a target and a

non-target for equal numbers of trials. Previous studies address this issue by sampling letters randomly. However, as discussed above, random sampling of individual trials is only reliable if the sample size is large. Experiments with a small number of trials may risk confounding the target factor with the letter factor. SweetPea can be used to address this problem directly, by balancing the letter factor with the derived target factor. Listing 3 shows SweetPea code that generates such an experiment sequence.

```

1 from sweetpea.primitives import factor, derived_level, window
2 from sweetpea import fully_cross_block, synthesize_trials_non_uniform, print_experiments
3
4 # letter factor
5
6 all_letters = ["A", "B", "C", "D", "E", "F"]
7
8 letter = factor("letter", all_letters)
9
10 # target factor
11
12 def is_target(letter):
13     return letter[0] == letter[2]
14 def is_no_target(letter):
15     return not is_target(letter)
16
17 target = factor("target", [
18     derived_level("yes", window(is_target, [letter], 3, 1)),
19     derived_level("no", window(is_no_target, [letter], 3, 1))
20 ])
21
22 # lure factor
23
24 def is_lure(letter, target):
25     return target[1] == "no" and letter[0] == letter[1]
26 def is_no_lure(letter, target):
27     return not is_lure(letter, target)
28
29 lure = factor("lure", [
30     derived_level("yes", window(is_lure, [letter, target], 2, 1)),
31     derived_level("no", window(is_no_lure, [letter, target], 2, 1))
32 ])
33
34 # experiment
35

```



```

36 design      = [letter, target, lure]
37 crossing    = [letter, target]
38 block       = fully_cross_block(design, crossing, [])
39
40 experiments  = synthesize_trials_non_uniform(block, 1)
41
42 print_experiments(block, experiments)

```

Listing 3: Example of a 2-Back experiment.

Lines 1-2 import relevant SweetPea modules, as described in the previous two sections. Note, however, that we also include `window` (line 1) to derive factor levels from sequences of more than two trials.

Lines 6-8 define the regular factor `letter` with its six levels.

The `target` factor is defined in lines 12-20. The predicates implementing each level of the factor `target` receive a list as input, similar to levels for transition factors described in the previous example. For instance,

```

def is_target(letter):
    return letter[0] == letter[2]

```

expects a list labeled `letter`. Each element of the list refers to a different trial within a specified window of trials relative to the current trial. In the function `is_target`, the argument coding for the factor `letter` is treated as a window of size 3. The last element (`letter[2]`) refers to the letter on the current trial and the first element (`letter[0]`) refers to the letter two trials back. A trial is considered a target if the letter on the current trial matches the letter two trials back. The window is specified in the declaration of the derived level `"yes"` for the `target` factor (line 18):

```

derived_level("yes", window(is_target, [letter], 3, 1)),

```

`window` is given the predicate `is_target` that returns true if the sequence satisfies requirements for this level. It passes the list of factors `[letter]` to the function `is_target`. The last two arguments of `window` define the window size—the number of past trials to consider, including the current trial—as well as the stride. The stride determines the number of trials to skip between the trials that are considered when selecting the new, derived level. For instance, if we were to determine the presence of a

target on every other trial (instead of every trial), we may specify a stride of 2 (instead of a stride of 1).

Lines 24-32 define the `lure` factor in an analogous manner, but using a window size of 2. Note that a trial is considered a lure if it is not a target and if the letter on the previous trial is the same as the letter on the current trial. The `is_lure` predicate, used to determine whether a trial is a lure, takes two arguments, `letter` and `target`,

```
def is_lure(letter, target):
    return target[1] == "no" and letter[0] == letter[1]
```

with both factors passed as a window of size two. Thus, `target[1]` and `letter[1]` refer to the target and letter factors, respectively, on the current trial and `letter[0]` refers to the letter factor on the previous trial. The window size of two trials for the predicate `is_lure` is specified in line 30:

```
derived_level("yes", window(is_lure, [letter, target], 2, 1))
```

Finally, line 36 lists the factors to consider in the design and line 37 defines the crossing between `letter` and `target`. The design and crossing are embedded without constraints in an experiment block (line 38). Line 40 generates an experiment sequence from the block with non-uniform sampling. Line 42 displays the output.

Table 3 shows an example output. Note that the target is not defined for the first two trials and the lure is not defined for the first trial due to a window size of two and one, respectively.

Discussion

The SweetPea language is intended to provide a format for describing experimental designs in a declarative form that is: (a) as concise and natural as possible; (b) sufficiently precise as to allow the application of standard computational algorithms for sampling and sample analysis; and (c) sufficiently general as to be useful for the widest possible range of experimental applications. Fully achieving all of these goals is of course a considerable challenge. Here, we present a first implementation, that we hope takes a meaningful step in this direction, by providing at least one concrete

formulation of how to approach these challenges, that is already useful for some applications, and that provides a foundation and benchmark for future development.

The experimental design challenges addressed by SweetPea are faced by researchers in other domains, such as neuroscience and machine learning. The measurement of neural correlates is often confounded by the order in which stimuli appear, especially if measurement techniques have a low temporal resolution (Aarabi, Osharina, & Wallois, 2017; Gorgolewski, Storkey, Bastin, Whittle, & Pernet, 2013). SweetPea addresses this problem, by allowing the user to counterbalance transition factors, and to impose sequential constraints on the generated sequence of trials. Another issue concerns proper sampling of training data for machine learning systems. Improper sampling of stimuli can bias statistical learning systems to misrepresent the space of all stimuli, leading to poor generalization performance (Zadrozny, 2004). Machine learning researchers can leverage SweetPea to specify the space of stimuli in terms of factors and levels, and use it to generate counterbalanced training data.

Despite the broad scope of experiments that can already be expressed in SweetPea, there are several ways in which the language can be improved. By making the code available in an open source repository (<https://github.com/sweetpea-org>), we hope that the community can help contribute to that effort. Toward that end, below we discuss potentially valuable directions for further development.

Continuous factors. SweetPea frames experimental design as a Boolean satisfiability problem, and thus requires that factors are composed of discrete levels. However, some experiments may require continuous factors, such as intertrial intervals or monetary rewards offered based on participant performance. At present, SweetPea can represent continuous factors by binning them into a limited set of discrete levels. However, a more efficient strategy would be to allow the user to sample these factors from pre-specified distributions (e.g. sampling the factor “intertrial interval” from an exponentially modified Gaussian distribution).

Uniform sampling of complex experimental designs. The current version of SweetPea can provide a SAT-sampler (Unigen; Chakraborty et al., 2014) with a general

logical formula describing the experiment suitable for use by automated sampling methods. Ideally, these would sample uniformly from the space of all possible solutions to the general logical formula provided by SweetPea. However, the current generation of SAT-samplers are not efficient enough to handle the scale of experimental designs that can be formulated using SweetPea. The uniform sampling of solutions to general logical formulas is still an area of active research in computer science. SweetPea stands to leverage progress made in this direction, by describing experimental designs in a form that can be used by SAT-samplers. In the future, this may make it possible to design larger experiments with the guarantee that all constraints are satisfied and solutions are uniformly sampled.

Debugging experimental designs. In some cases, a researcher may over-constrain an experimental design, such that there exists no possible experiment that satisfies all constraints. In these cases, SweetPea can inform the user that the experiment is overspecified. However, at present, SweetPea can only offer limited insight into which constraints are in conflict with one another. Therefore, a useful target for ongoing development of SweetPea would be tools that aid researchers in debugging their experimental design. This could involve iterative procedures that achieve partial satisfiability by dropping design constraints with low priority.

Automatic design specification. Another useful avenue for development would be to automatically derive the minimal set of counterbalancing conditions that satisfies the statistical analysis specified by the researcher. A major goal of experimental design is to warrant proper statistical inference from the data generated by the experiment. As a consequence, counterbalancing schemes and sequential constraints are often determined by the statistical analysis that a researcher attempts to perform. For instance, a researcher who seeks to contrast two levels of an experiment factor with a Student's t-test, would want to counterbalance the two levels with respect to all levels of nuisance factors. This can be achieved by distributing the levels of nuisance factors uniformly across the levels of factors of interest. SweetPea could assist in generating such a counterbalancing scheme, but currently the specification itself would have to be

performed by the researcher. More generally, a particularly promising direction for future work would be the integration of a high-level interface for specifying an experimental design based on a desired statistical analysis (e.g. contrasts defined by a Student's t-test, ANOVA or linear mixed model).

We plan to continue work along these lines, and hope others will join us in this effort. Doing so offers the promise of providing a rigorous and standardized method for experimental design that would help avert many of the problems that have been identified in past and current experimental research.

Acknowledgements of Support

This project was made possible through the support of a grant from the John Templeton Foundation, as well as the National Science Foundation (NSF-1813123). The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation. We thank Markus Spitzer for providing helpful feedback on the tutorial.

References

- Aarabi, A., Osharina, V., & Wallois, F. (2017). Effect of confounding variables on hemodynamic response function estimation using averaging and deconvolution analysis: an event-related nirs study. *Neuroimage*, 155, 25–49.
- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., . . . others (2016). Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)* (pp. 265–283).
- Allport, A., & Wylie, G. (1999). Task-switching: Positive and negative priming of task-set.
- Bergstra, J., & Bengio, Y. (2012). Random search for hyper-parameter optimization. *The Journal of Machine Learning Research*, 13(1), 281–305.
- Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5), 365.
- Chakraborty, S., Meel, K. S., & Vardi, M. Y. (2014). Balancing scalability and uniformity in sat witness generator. In *2014 51st acm/edac/ieee design automation conference (dac)* (pp. 1–6).
- Chapelle, O., Vapnik, V., Bousquet, O., & Mukherjee, S. (2002). Choosing multiple parameters for support vector machines. *Machine learning*, 46(1-3), 131–159.
- Cherkaev, A. (2019). *SweetPea: a language for experimental design* (Unpublished master’s thesis). The University of Utah, Salt Lake City, UT, USA.
- Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., & Smith, E. E. (1997). Temporal dynamics of brain activation during a working memory task. *Nature*, 386(6625), 604–608.
- Collaboration, O. S., et al. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716.
- Cooper, S., & Marí-Beffa, P. (2008). The role of response repetition in task switching. *Journal of Experimental Psychology: Human Perception and Performance*, 34(5),

1198.

- Dale, A. M. (1999). Optimal experimental design for event-related fmri. *Human brain mapping*, 8(2-3), 109–114.
- De Leeuw, J. R. (2015). jspsych: A javascript library for creating behavioral experiments in a web browser. *Behavior research methods*, 47(1), 1–12.
- Dijksterhuis, A., Van Knippenberg, A., & Holland, R. W. (2014). Evaluating behavior priming research: Three observations and a recommendation. *Social Cognition*, 32(Supplement), 196–208.
- Druery, C., & Bateson, W. (1901). Experiments in plant hybridization. *Journal of the Royal Horticultural Society*, 26, 1–32.
- Drummond, C. (2006). Machine learning as an experimental science (revisited). In *Aaai workshop on evaluation methods for machine learning* (pp. 1–5).
- Gardner, M., Neumann, M., Grus, J., & Lourie, N. (2018, 10). Writing Code for NLP Research. In *Proceedings of the 2018 conference on empirical methods in natural language processing: Tutorial abstracts*. Melbourne, Australia: Association for Computational Linguistics.
- Gorgolewski, K. J., Storkey, A. J., Bastin, M. E., Whittle, I., & Pernet, C. (2013). Single subject fmri test–retest reliability metrics and confounding factors. *Neuroimage*, 69, 231–243.
- Gureckis, T. M., Martin, J., McDonnell, J., Rich, A. S., Markant, D., Coenen, A., . . . Chan, P. (2016). psiturk: An open-source framework for conducting replicable behavioral experiments online. *Behavior research methods*, 48(3), 829–842.
- Hartshorne, J. K., de Leeuw, J. R., Goodman, N. D., Jennings, M., & O’Donnell, T. J. (2019). A thousand studies for the price of one: Accelerating psychological science with pushkin. *Behavior research methods*, 1–22.
- Jou, J. (2014). Task-switching cost and repetition priming: Two overlooked confounds in the fixed-set procedure of the sternberg paradigm and how they affect memory set-size effects. *Quarterly Journal of Experimental Psychology*, 67(10), 1871–1894.
- Kiesel, A., Steinhauser, M., Wendt, M., Falkenstein, M., Jost, K., Philipp, A. M., &

- Koch, I. (2010). Control and interference in task switching—a review. *Psychological bulletin*, 136(5), 849.
- Klein, R. A., Ratliff, K. A., Vianello, M., Adams Jr, R. B., Bahník, Š., Bernstein, M. J., ... others (2014). Investigating variation in replicability. *Social psychology*.
- Krause, F., & Lindemann, O. (2014). Expyriment: A python library for cognitive and neuroscientific experiments. *Behavior Research Methods*, 46(2), 416–428.
- Kühberger, A., Fritz, A., & Scherndl, T. (2014). Publication bias in psychology: a diagnosis based on the correlation between effect size and sample size. *PloS one*, 9(9), e105825.
- Langley, P. (1988). Machine Learning as an Experimental Science. *Machine Learning*, 3(1), 5–8.
- Logan, G. D., & Schneider, D. W. (2010). Distinguishing reconfiguration and compound-cue retrieval in task switching. *Psychologica Belgica*, 50(3), 413–433.
- Mayr, U., & Keele, S. W. (2000). Changing internal constraints on action: The role of backward inhibition. *Journal of Experimental Psychology: General*, 129(1), 4.
- Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6), 1423.
- Mendel, G. (1866). Versuche uber pflanzen-hybriden. *Verhandlungen des naturforschenden Vereins in Brunn fur*, 4, 3–47.
- Merkel, D. (2014). Docker: lightweight linux containers for consistent development and deployment. *Linux journal*, 2014(239), 2.
- Miłkowski, M., Hensel, W. M., & Hohol, M. (2018). Replicability or reproducibility? on the replication crisis in computational neuroscience and sharing only relevant detail. *Journal of computational neuroscience*, 45(3), 163–172.
- Miller, G. (2011). *The mating mind: How sexual choice shaped the evolution of human nature*. Anchor.
- Myung, J. I., & Pitt, M. A. (2009). Optimal experimental design for model discrimination. *Psychological review*, 116(3), 499.

- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., . . . others (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems* (pp. 8026–8037).
- Peirce, J. W. (2007). Psychopy—psychophysics software in python. *Journal of neuroscience methods*, 162(1-2), 8–13.
- Peirce, J. W. (2009). Generating stimuli for neuroscience using psychopy. *Frontiers in neuroinformatics*, 2, 10.
- Peng, R. D. (2011). Reproducible research in computational science. *Science*, 334(6060), 1226–1227.
- Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of experimental psychology: General*, 124(2), 207.
- Rossi, J. S. (1990). Statistical power of psychological research: What have we gained in 20 years? *Journal of consulting and clinical psychology*, 58(5), 646.
- Sherman, R., & Pashler, H. (2019). Powerful moderator variables in behavioral science? don't bet on them (version 3).
- Sochat, V. V., Eisenberg, I. W., Enkavi, A. Z., Li, J., Bissett, P. G., & Poldrack, R. A. (2016). The experiment factory: standardizing behavioral experiments. *Frontiers in psychology*, 7, 610.
- Stroebe, W., & Strack, F. (2014). The alleged crisis and the illusion of exact replication. *Perspectives on Psychological Science*, 9(1), 59–71.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of experimental psychology*, 18(6), 643.
- Sudevan, P., & Taylor, D. A. (1987). The cuing and priming of cognitive operations. *Journal of Experimental Psychology: Human perception and performance*, 13(1), 89.
- Wells, G. L., & Windschitl, P. D. (1999). Stimulus sampling and social psychological experimentation. *Personality and Social Psychology Bulletin*, 25(9), 1115–1125.
- Zadrozny, B. (2004). Learning and evaluating classifiers under sample selection bias. In *Proceedings of the twenty-first international conference on machine learning*

(p. 114).

Zmigrod, S., & Hommel, B. (2013). Feature integration across multimodal perception and action: a review. *Multisensory research*, 26(1-2), 143–157.

Trial	color	word	response	congruency
1	brown	red	down	incongruent
2	green	red	right	incongruent
3	brown	blue	down	incongruent
4	red	blue	left	incongruent
5	brown	green	down	incongruent
6	blue	red	up	incongruent
7	red	green	left	incongruent
8	green	brown	right	incongruent
9	red	brown	left	incongruent
10	green	blue	right	incongruent
11	blue	green	up	incongruent
12	blue	brown	up	incongruent
13	blue	brown	up	incongruent
14	green	red	right	incongruent
15	brown	blue	down	incongruent
16	red	brown	left	incongruent
17	brown	green	down	incongruent
18	blue	red	up	incongruent
19	red	green	left	incongruent
20	green	brown	right	incongruent

Table 1

Example solution to a Stroop color naming design.

Trial	color	word	task	congruency	response	task_transition	response_transition
1	red	green	color naming	incongruent	left		
2	red	red	word reading	congruent	right	switch	switch
3	green	green	color naming	congruent	right	switch	repetition
4	green	red	color naming	incongruent	right	repetition	repetition
5	green	green	word reading	congruent	left	switch	switch
6	green	green	word reading	congruent	left	repetition	repetition
7	red	green	color naming	incongruent	left	switch	repetition
8	red	red	color naming	congruent	left	repetition	repetition
9	green	red	color naming	incongruent	right	repetition	switch
10	red	red	word reading	congruent	right	switch	repetition
11	red	green	word reading	incongruent	left	repetition	switch
12	green	red	word reading	incongruent	right	repetition	switch
13	green	green	word reading	congruent	left	repetition	switch
14	red	red	color naming	congruent	left	switch	repetition
15	red	green	word reading	incongruent	left	switch	repetition
16	green	red	color naming	incongruent	right	switch	switch
17	red	red	color naming	congruent	left	repetition	switch
18	green	red	word reading	incongruent	right	switch	switch
19	green	red	color naming	incongruent	right	switch	repetition
20	green	red	word reading	incongruent	right	switch	repetition
21	green	red	word reading	incongruent	right	repetition	repetition
22	red	red	word reading	congruent	right	repetition	repetition
23	red	green	color naming	incongruent	left	switch	switch
24	red	green	color naming	incongruent	left	repetition	repetition
25	green	green	word reading	congruent	left	switch	repetition
26	red	green	word reading	incongruent	left	repetition	repetition
27	red	red	word reading	congruent	right	repetition	switch
28	red	red	color naming	congruent	left	switch	switch
29	green	green	color naming	congruent	right	repetition	switch
30	green	green	color naming	congruent	right	repetition	repetition
31	red	green	word reading	incongruent	left	switch	switch
32	green	green	color naming	congruent	right	switch	switch
33	red	green	color naming	incongruent	left	repetition	switch

Table 2

Example solution to a task switching design.

Trial	letter	target	lure
1	A		
2	B		no
3	B	no	yes
4	D	no	no
5	B	yes	no
6	D	yes	no
7	F	no	no
8	C	no	no
9	F	yes	no
10	C	yes	no
11	A	no	no
12	E	no	no
13	A	yes	no
14	E	yes	no

Table 3

Example solution to a 2-Back task design.