Implementing an Open-access, Data Science Programming Environment for Learners

Austin Cory Bart, Javier Tibau, Eli Tilevich, Clifford A. Shaffer, Dennis Kafura
Computer Science
Virginia Tech
Blacksburg, Virginia, USA
acbart@vt.edu, jtibau@vt.edu, tilevich@vt.edu, shaffer@vt.edu, kafura@vt.edu

Abstract—A key retention issue when educating computing novices is ensuring that the frustrations of mastering programming fundamentals do not demotivate and discourage students from studying the discipline. In particular, non-CS majors often struggle to find relevance in traditional computing curricula that tend to either emphasize abstract concepts, focus on non-practical entertainment (e.g., game and animation design), or rely on decontextualized settings. To address these issues, this paper introduces BlockPy, a block-based environment for Python (http://www.blockpy.com). BlockPy is a web-based, open-access programming environment that supports introductory programming with an emphasis on data science. It promotes long-term transfer by scaffolding an introduction to textual programming (Python) through a block-based programming view, ideal for beginners of any background. By supporting the latest Learning Tools Interoperability (LTI) standards, BlockPy is designed to support both informal learners and formal class settings. Specifically, it can be configured to provide guiding feedback for its interactive programming problems, so as to support learners at their own pace. The results from a pilot study of the initial deployment and utilization of BlockPy indicate the potential of the environment to address many of the problems faced by novice learners.

Keywords-Computer science education; Computer aided instruction; Data analysis; Web services;

As computing becomes pervasive in our society across fields, working professionals increasingly need to acquire some expertise in computing in addition to their core domain knowledge. General education computing curricula at the university level (e.g., "Computational Thinking" courses) are scaling, Massively Open Online Courses are flourishing, and a large class of learners are pursuing ad-hoc, non-formal learning experiences on their own. Both these traditional and non-traditional learners often have little experience with computing, low self-efficacy, and are uncertain about the value of computing towards their long-term career goals. Not only do they need special scaffolding unique to their ability and motivational level, but they also need fewer barriers in the technology they access these materials with. Our solution to serve this population is BlockPy: an open-access, web-based Python environment for data

science that supports learners with guided instruction and an accessible interface (http://www.blockpy.com).

Why Data Science? Modern approaches to contextualizing introductory courses have focused on making the experience "fun" and "interesting", with an emphasis on game design and media computation [1]. However, student motivation is a complex, multi-faceted construct dependent on more than just situational interest; in particular, holistic models of motivation suggest that students also need to feel that the material is inherently useful to learn, and that long-term career goals are being satisfied [2]. Despite an attempt at convincing students otherwise, Media Computation is not perceived as an authentically useful context for non-majors, based on a study by Guzdial et al [3].

We suggest that Data Science is a motivating context that can appeal in a different way to students, thanks to the wide-spread need for data processing in other majors. Students are often studying computing to learn how to manage the dizzying quantities of data being stored, used, and analyzed in a discipline or for a specific self-derived project. By grounding the content in this context, students can be more easily convinced of the relevance of computing and understand how the materials fit together more clearly. By aligning the context with students' long-term needs, students can also learn skills more relevant to their needs. Finally, data science as a context naturally lends itself to teaching topics related to structured data, iteration, and other core material, making it a pedagogically valuable context to the CS instructor.

Why Python? Python has become one of the most popular introductory programming languages [4], thanks to its simple syntax but impressive power, including strong support for data science thanks to popular libraries like MatPlotLib. Python requires little code in order to accomplish interesting things, so novices are not bogged down with a tremendous amount of syntactical details. Its wide-spread use in both introductory classes and industry motivated our choice.

Why Blocks? However, any kind of programming is

still a challenge to beginners, due to the nature of coding as the "most powerful, but least usable human-computer interface ever invented." [5] Block-based languages have been shown to mitigate the start-up time for students to start programming and accomplishing tasks [6], [7]. By providing structure and an immediate view into the entire user interface of a language, blocks greatly benefit introductory learners. Echoing a popular sentiment among CS educators, *"At the novice programming level, blocks-based languages are the most promising direction today"* [8].

Why Another Python Web Environment? There are several environments available today that let students and instructors write Python in the browser, including CodeSkulptor [9], Pythy [10], and the Online Python Tutor [11]. BlockPy stands on the shoulders of giants, integrating features inspired by these environments and introducing novel ones. But none of these existing Python environments scaffolds transitioning students into textual programming languages.

BlockPy was designed to provide dual support for both block-based and text-based code authoring. At any time, the student can switch freely between a block-based view of their code and a traditional text-based view. This powerful feature is inspired by Pencil Code, which uses its own Logo language [12], and similar implementations have been successful as a fading scaffold for students [13].

BlockPy extends Pythy's [10] support for "assignments", which are problems that integrate presentation with assessment. However, Pythy only supports traditional unit testing to provide students with feedback, while BlockPy provides an API for code analysis and free-form text guidance that instructors can configure to give helpful suggestions to their students. Furthermore, Pythy has limited support for data science, whereas BlockPy has a rich library of data sources to draw on and a MatPlotLib-based plotting API.

CodeSkulptor, Pythy, and BlockPy all use the same internal engine for running Python code ("Skulpt"). Although CodeSkulptor has an extensive API for creating user interfaces and games, it suffers from using a non-standard library. Although suitable for beginners, this library does not aid the transition to textual programming languages. In BlockPy, the philosophy is to maintain compatibility with existing systems when possible. Instead of a custom plotting API, for instance, we mimic the MatPlotLib interface.

Finally, OnlinePythonTutor has proven to be an incredibly useful tool for visualizing program state in Python programs. However, we hypothesize that the depth of detail that it gives can be overwhelming for introductory students (e.g., the visualizations use terminology such as Frames and Objects, which might be foreign to students despite being proper terminology). BlockPy's state explorer does not attempt to match Python Tutor's thoroughness or accuracy, but instead is targeted at providing as helpful as picture of program state as possible within the constraint of simplicity. This includes highlighting data types and console output as part of the visualizers' program state, as shown in Figure 1 on the right side. Additionally, we remove Python-Tutor's server-based python dependency by relying on Skulpt, which works entirely in the students' browser.

I. BLOCKPY

The primary design goals of BlockPy are as follows.

1) Reduce the barriers to learning programming.
2) Promote authenticity by empowering students to complete real-world problems.
3) Promote maturity by faded scaffolds (e.g., transitioning from blocks to text).
4) Minimize the need for help from human instructors, except where truly needed.

A. Open-source, Open-access

The fundamental vision of BlockPy is a highly accessible, web-based platform for anyone to learn how to program. All of its code is completely open-source, and it leverages a number of open-source libraries. There is no registration barrier to start using the software. Although there are a few features that benefit from registration, such as managing classrooms, registration is free. Guided learning materials developed in the system are meant to be shared within the ecosystem of educators who use BlockPy.

As an editor, BlockPy constantly stores user code as it is developed. These logs are stored at the keystroke level to aid future program analysis (described further in section I-E). The latest version of the user's code is therefore available between sessions. When operating in offline mode, the code is stored in the `LocalStorage` browser object; when and if the connection is reestablished, synchronization is performed. The system has special logic to ensure that only the most recent version of the author's code is stored as the canonical reference; this is crucial for situations where BlockPy is used in formal classroom environments.

B. Python Execution

The BlockPy system is built to work offline first, ideal for places where internet connectivity is neither assured nor reliable. Python Code Execution is achieved through a modified instance of the Skulpt JavaScript library. Skulpt is a full Python parser and compiler, supporting almost all of the language features by generating runnable JavaScript code. However, it only has partial support for the complete Python standard library,
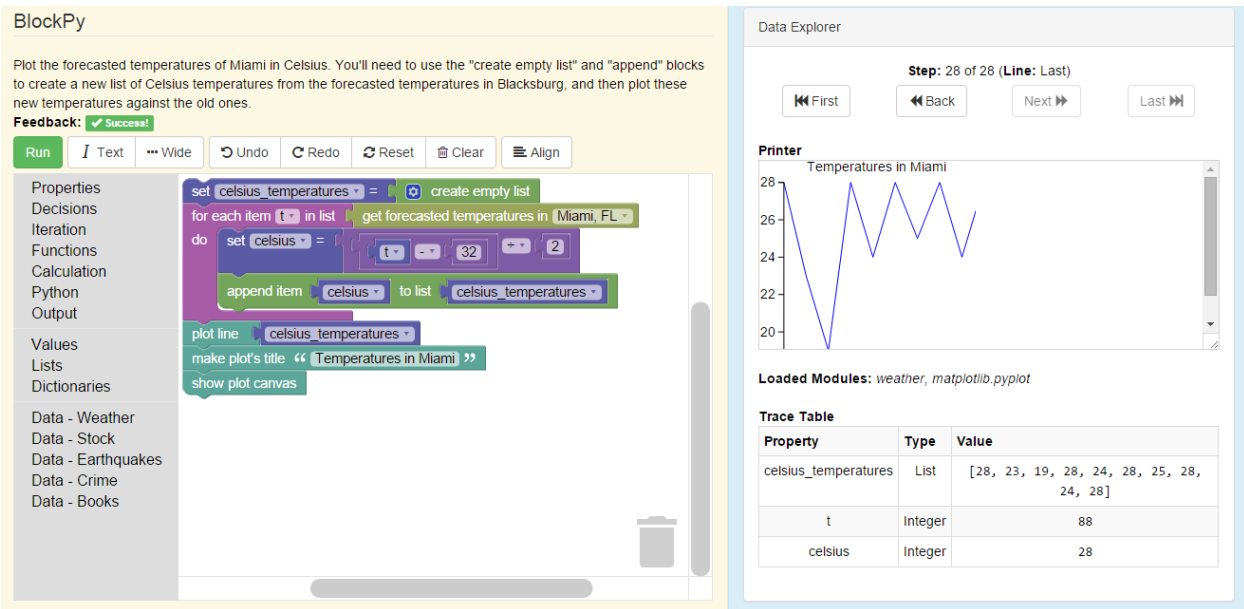
Figure 1. A complete screenshot of BlockPy in action

which includes a tremendous number of libraries. The Skulpt execution environment resides entirely within the users' browser, so there is no reliance on an external server except for the initial page-load. Libraries that depend on external server data can operate in an offline-mode, making use of cached data. The long-term goal of this project is to support a set of rich libraries so that sophisticated applications can be developed beyond console-based problems.

C. Block-based Python

To support introductory learners as they grapple with Python syntax, the initial interface in BlockPy is block-based, using a modified version of the popular Blockly JavaScript library. Language features (iteration, decision, variable assignment and access, etc.) are contained in a toolbox on the left-side, from which users drag-and-drop blocks onto a canvas. At any time, BlockPy's block interface can only generate syntactically valid Python code, enforced by the "snapping" connectors of the blocks (although it is possible to generate semantically incorrect code – this is further discussed later). Users are given visual refactoring tools to organize the workspace and perform typical IDE commands (e.g., copy, paste, clear). An additional advantage of the Block-based interface is that blocks can be manipulated using touch interfaces, particular on mobile devices. This block interface is also supported by a text interface; the interaction of these two interfaces is described more fully in section I-H.

An important question is how many language details should be exposed, and at what rate. A rarely used feature of `for` loops in Python is to contain an `else` clause that is executed upon successful completion of the loop (that is, when it is not prematurely escaped using a `break` statement). This advanced language feature is meant to draw special attention to connected actions that must be performed after the iteration is completed (similar to a `finally` statement with exceptions). However, if an `else` clause were made available to beginners first trying to grapple with iteration, it is likely they would confuse the concept with the conditional `else` clause used in `if` statements. Cognitive Load Theory can be a harsh mistress for beginners, and the user interface needs to avoid exposing unnecessary details where possible. While hiding `else` bodies in a `for` loop is a clear case, there more subtle examples. It can be very difficult to recognize when the learner is ready to use parallel assignment (which can be useful for splitting strings with a known format, such as "month-year"), and therefore should be able to specify multiple variables on the left side of an assignment block. A block-based language forces a teacher to make important decisions about how to expose language features. As part of the future work of BlockPy, we're experimenting with exposing language features at different rates, adjustable by the instructor or adaptable to a learners' rate. In theory, the system can expose a more accurate language model, similar to what happens in other environments such as Dr. Racket [14].

D. Adaptive Guided Practice

One of the most powerful features of BlockPy is the interactive, guided feedback feature. A limitation of

programming environments like Snap! is that they are not pedagogically interactive – students completing an assignment in the system are not guided to success. It is up to the learner to decide when they have completed their program, and whether it is both correct and meets the specification. For independent learners outside of a formal learning experience, this makes progressing extremely difficult and requires high levels of self-regulation and metacognition. The adaptive elements follow an Expert model; when students run their code, it is checked against instructor-provided logic and constraints. If the student code fails for some reason, they can be offered a suggestion or motivational remark. Correct code gives a green checkmark through the system – we have found that the motivating power of a positive marker cannot be overestimated.

In the Instructor Mode (shown in Figure I-K), teachers and developers can provide a problem instance. First, a WYSIWYG rich text editor edits the problem description, supporting images and any valid HTML content. Second, the instructor can provide code in several special canvases that affect the students experience. Instructors write this code using the same text/block interface that students use, and can quickly switch out of Instructor Mode to test their creation. The first canvas is the "default code", shown to the student when they first begin the problem so that they are not facing a blank canvas. The second canvas defines a special `on_run` function. This function is the interface for providing feedback and consumes the students code, their final output, and a complete trace of their programs state. The function returns either a boolean indicating success, or an HTML string that is rendered to the user as feedback. The instructor is free to write whatever logic they want, such as searching for a specific line of code, testing the outputs on the console, or walking through the programs state to satisfy certain invariants. An API for common checks is evolving based on common use cases, such as checking whether a student is iterating over a non-empty list (a surprisingly common problem) or parsing the program's Abstract Syntax Tree to ensure that they have properly used a specific language feature. We are also exploring what kinds of interventions an instructor can make beyond rendering textual feedback; perhaps displaying a pop-up dialog with an embedded instructional video, or alerting an instructor to provide Just-In-Time instruction for a particular struggling case.

E. Program Analysis for Deeper Learning

We have recently started experimenting with simple program analysis techniques to find both general mistakes that novices make, and problem-specific errors. For example, our observation is that beginners often fail to understand the true purpose of certain variables, and
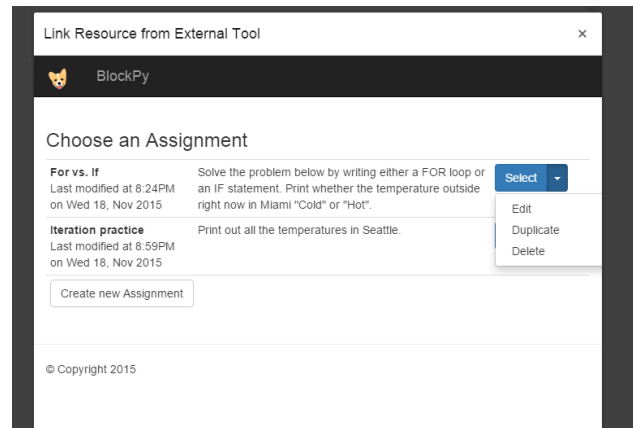

Figure 2.   Teacher Perspective for Managing BlockPy Assignments

include them in a mimicked style. By performing simple variable liveness analyses, we can identify these variables and raise a serious warning. Most modern editors feature this kind of analysis, but usually trust that the user has a reason and behaves non-invasively; we can make stronger statements about our users' ability levels in order to make stronger recommendations against such cargo-cult programming.

As previously mentioned, we are also securely recording the entire log of users' programs and the interaction log with the system. Our hope is that by mining this repository of code, we can infer common patterns that suggest undesirable learner behavior. For example, users that frequently move the same blocks without progressing in the problem objectives ("churning" the interface) might be indicative of taking longer on the problem than other users. Alternatively, students who pick a decision block to complete a problem about iteration might commonly fail to complete a problem. By proving or disproving such hypotheses, we can improve the automatic feedback of the system and provide more individualized support predicts future student success based on current behavior.

F. LTI Support

A growing cause for alarm in the education community is the ever-expanding array of third-party tools that demand control over students data. The BlockPy project is committed to supporting LTI technology to reduce instructors dependency on BlockPy for course management. LTI (Learning Technology Interoperability) is a mechanism by which instructors can embed questions in their existing course management software (e.g., Canvas, OpenEdX) and receive assignment outcomes such as performance and participation data.

Currently, BlockPy has a dual-layer registration model. A typical learner can register in the system

without ever being aware of the LTI technology. However, an instructor using the system can obtain a secret key and configuration URL that is used within their Learning Management System. After BlockPy executes the LTI dance with the LMS, students on the course website may use Blockpy without registering for an account – the first time they log in through their LMS' provided link, they are invisibly registered in the system with a regular account supplemented with additional information from canvas (unless an existing account can be found for the LMS or for the unique information provided by the system). As students complete work, assignment progress is reported to the LMS so that they can receive a grade (additional user profile and code logs are stored on the BlockPy servers, although we will offer support for instructors to turn this off as needed for privacy purposes). Instructors can use a special interactive menu for managing exercises associated with a course, as demonstrated in figure 2. Currently, the system supports LTI version 1.1, but as LTI version 2.0 rolls out in the coming years, we intend to support the latest standard.

G. Data Science Blocks

BlockPy focuses on Data Science as its primary context, and so we have specific blocks and APIs for working with data. We use an existing fork of Skulpt with the popular MatPlotlib library, and extended it with our own implementation of the CORGIS project [15]. The MatPlotLib API, for example, provides a "plot(list)" function to create simple line plots. By basing everything around MatPlotLib, students can seamlessly shift to a serious Python programming environment without loss of code or productivity.

The CORGIS (Collection of Real-time, Giant, Interesting, Situated datasets) Project seeks to make motivating datasets available to introductory students through simple programming libraries [15]. For instance, up-to-date temperature and humidity reports for American cities are available through the Weather library, These datasets are drawn from many disciplines, resulting in material meant to be universally interesting and relevant. These datasets are also equipped with scaffolding that simplifies the process of working with what are sometimes complicated data sources; for instance, real-time data sources are cached during development so that students are not reliant on a consistent internet connection. Currently, BlockPy supports five different CORGIS libraries: weather data, stock trading data, earthquake data, United States crime reports, and a small book dataset. The newest version of CORGIS is being extended to export libraries that are directly usable via BlockPy.
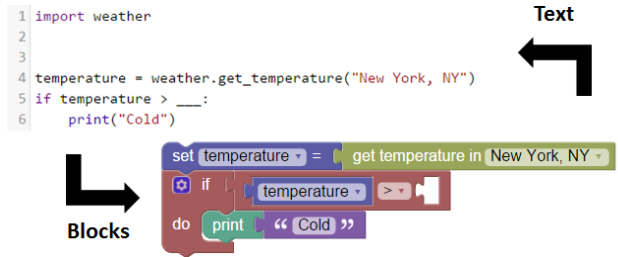

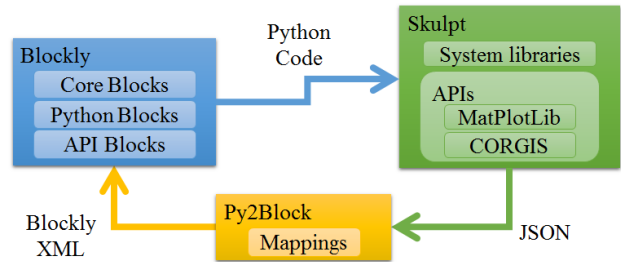Figure 3.   User perspective of block/text transition

Figure 4.   Architectural perspective of the block/text transition

H. Mutual Language Translation

One of the technical contributions of this project is the Mutual Language Translation between Blockly and Python. Figure 4 gives an overview of the technical architecture used to achieve this translation. Blockly outputs valid python source code, which can be passed into Skulpt in order to extract a JSON representation of the Abstract Syntax Tree. Alternatively, the Python source of the Skulpt program can be edited directly in the Text view. Either way, this AST is parsed using our Py2Block library to generate an XML representation that Blockly can render in the Block View. Figure 3 demonstrates the users' experience.

Blockly already supports compilation of its blocks to Python, JavaScript, PHP, and Dart. However, this multiple language support comes at a cost of reduced isomorphism—each language has different syntax for their common operations, and it is impossible to create a fully-featured block language with a one-to-one mapping between them. For example, JavaScript has no support for parallel assignment, a commonly-used feature in Python, while Python does not have a unary increment operator. Blockly itself has syntax and vocabulary descended from Logo.

Instead of trying to satisfy multiple languages, we have dropped support for JavaScript and the rest in favor of a more full-featured mapping to Python. This requires minor changes that introduce Python-centric syntax details: function blocks are labeled "define", assignment blocks have an "=" symbol, the "add item to list" block is renamed to "append". Blockly has also
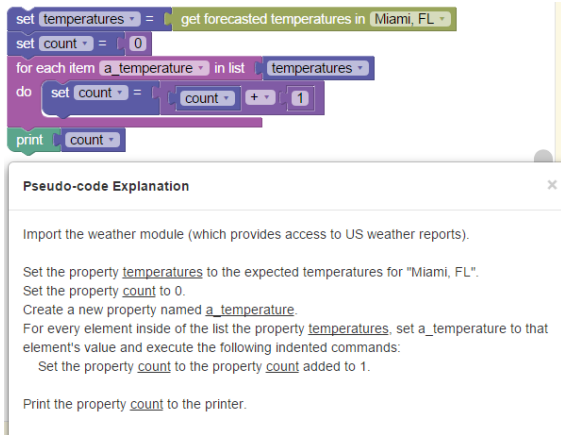
Figure 5. A Python program can also be rendered in Natural Language.

Figure 6. The State Explorer displays information about the entire state of the program, including output and loaded modules.

been extended with new language features, including dictionary access and creation.

Eventually, the interface should offer a complete isomorphic mapping to Python. However, there are a number of complications to resolve before that can occur. For instance, Python uses square brackets for both list indexing and dictionary access. There is a strong desire to differentiate between these types of access, visible in the block view as two distinct kinds of blocks ("get ith element of list" vs. "get key from dict" blocks). However, depending on what features of python are supported, it can be difficult or even impossible to statically identify the usage of a given pair of brackets– sophisticated program analysis techniques are needed.

I. Natural Language Code Explanation

Students can often struggle to interpret structured text into human language, because they are not familiar with the conventions implied by the unfamiliar syntax. A new feature of BlockPy is to generate natural language representations of programs using the same genera- tive technology that exports Python representations of blocks. This description, while just as regular as a Python program, expands on complex code elements to break down what they are doing at any step. Figure 5 demonstrates a potential natural language description for a program that measures the length of a list. There are a number of user interface optimizations that we are exploring – using tooltips to explain an operation in more detail, or juxtapositioning the natural language with the original Python and drawing comparisons be- tween the two.

J. Parson's Problems

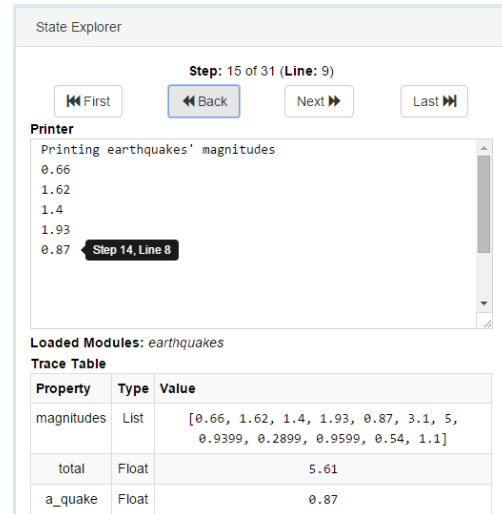Parsons' Problems are a special type of coding exer- cise where all of the necessary code blocks are present, but disconnected and shuffled. These kinds of problems scaffold beginners by giving them everything they need to complete the problem up front, reducing many of the barriers to getting started [16]. BlockPy supports these types of problems with a special "Parsons Mode" where top-level blocks are constantly shuffled in the block-mode. When a student tries to convert code with disconnected blocks, the generated Python code will be filled in with triple underscores, as demonstrated in Figure 3. These underscores (usually valid syntax in Python) will trigger a runtime error, encouraging students to think critically about their code.

K. State Explorer

A common debugging tool in many modern IDEs is a Variable Explorer, used to trace the programs' execution over time. BlockPy expands this concept into a State Explorer. The State Explorer displays more than just information about variables - the dashboard also reveals information about the programs' state over time. Users can step through the code's execution, affecting what is currently printed (including both console printing and the creation of graphs), imported modules, and the values and types of the variables. The relevant blocks and text are also highlighted appropriately.

II. Model Use Cases

In this section, we consider some example scenarios that describe our vision of typical BlockPy use cases. Our intent is for BlockPy to be useful in both formal and informal situations, with different features being especially helpful.
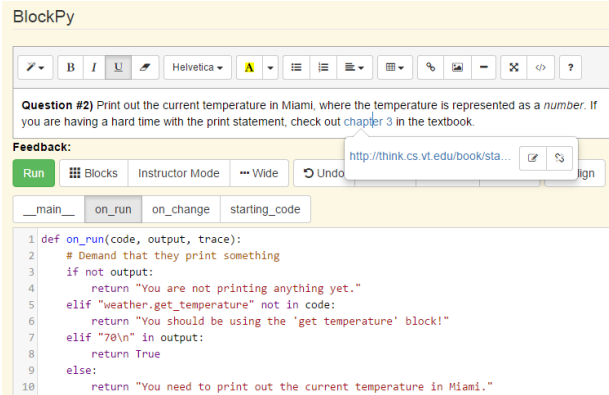
Figure 7. Instructor Mode lets you edit questions and and define special functions to give feedback.

A. Independent Learner

A learner independently logs into the BlockPy system and selects an introductory problem on calculating averages using iteration: *"Is the weather in Seattle above 60 degrees fahrenheight? Print Yes or No."* As a complete novice, they are unsure what to do after reading the problem description. If they decide to try and cheat by checking the current weather in Seattle and printing the literal value, the system intelligently notices that they are missing any use of a relevant weather block, and explains that they need to combine programmatic decision logic with the appropriate data source. They think to access the "Weather" block category, and grab the `get weather in city` block, but are unsure what to do next. When they run their program, the system notices that they have not used any decision constructrs (ie there is no `if` statement), and suggests reading a linked chapter in an online textbook about how to use it. If they continue to struggle with integrating pieces, the system can provide increasingly detailed hints until they succeed.

B. Classroom Lesson

Another common use case for the system would be an instructor with a large classroom of students. The instructor is using Canvas, an LTI-capable Learning Management System of growing popularity in higher education. They create a series of assignments for the day's classwork using the interface shown in figure 2. Students log into Canvas and begin working on the assignments. As they complete the assignments, their grade is reported to Canvas. The instructor can monitor progress for the class and check which students are struggling to complete assignments. This information can allow them to target under-performers earlier with interventions. The more automatic feedback that instruc-

tors make available, the less they need to focus on simple problems ("You were checking the temperature for the wrong city.") and the more they can focus on students that are truly struggling ("What is iteration?").

C. 1-1 Tutoring

On several occasions, we have found BlockPy to be a useful tool for correcting individual students' misconceptions. In particular, the block representation of programs' can help beginners grasp that code is not a series of symbols but a structured representation of an algorithm. Consider a student struggling to write the necessary syntax for indexing a nested dictionary (e.g., a crime report broken into multiple levels, with the burglary rate for a city nested under a violent crime categorization; this is an actual representation of how crime data is stored in our example datasets). The student may not have a clear image of how the layered structure of data can translate into a chain of dictionary accesses. By sitting with the student, the instructor could build up an expression accessing the data by connecting together dictionary access blocks to the data block, showing the generated Python code at each step. The learner can visually see how discrete chunks of the code correlate to discrete blocks, helping them overcome the beginner misconception of code as a random mishmash of characters.

III. PILOT STUDY

BlockPy has been piloted in an introductory "Computational Thinking" course with 35 students in Spring 2015. These students come from a diverse range of majors, including liberal arts (57%), architecture (17%), sciences (15%), and others (11%). There were 20 female students (57%) and 15 male. The vast majority of students reported no prior experience with programming or Computer Science, with less than 17% having taken the high school AP course. They were evenly distributed across years, with slightly more senior (29%), equal percentages of sophomore and juniors (26%), and fewer freshmen (14%). Although small, the student demographics reflect a varied population.

The Computational Thinking course's content is focused on teaching Abstraction and Algorithms. While programming is not a primary learning objective, it is an important topic in the course as a tool of concretely talking about the higher level objectives. The first third of the course, students work with NetLogo (although they do not program in it, they only read code) and participate in a number of explanatory kinesthetic activities. Then they are introduced to Python using BlockPy, for which they spend roughly six classes on completing guided practice problems in a mastery style (that is, they are allowed an infinite number of retries). The
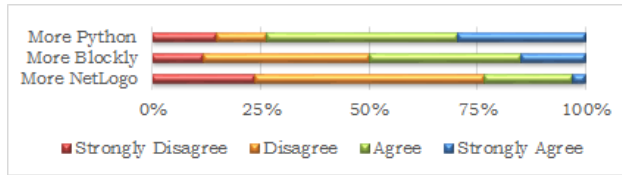
Figure 8. Do students want more time with NetLogo, BlockPy (Blockly), or Spyder (Python).
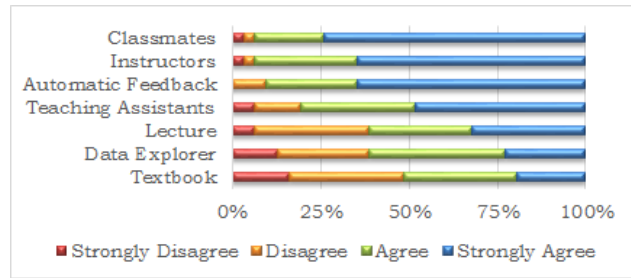

Figure 9. How helpful to students' learning were these resources?

Figure 10. How useful to students' long-term career goals was each of the following experiences?

next two classes are devoted to using a regular Python environment (Spyder) to complete small programming assignments (similar to the ones done with BlockPy). Finally, students are given eight class periods to work on their individual final project in Spyder.

A. Methodology

Every student interaction was logged by the system, including changes to the students' code. This data was also correlated to two surveys, one given after the BlockPy section and the other given at the end of the course. The survey was composed of Likert questions on a 4-point scale and open-ended qualitative questions. This gives us a wealth of data to analyze, although the small population size makes it difficult to derive significant results.

B. Perceptions of BlockPy

The first survey question students were asked was about whether they wanted more time with each of the programming environments they used in the course: NetLogo, BlockPy, or Spyder. The results of this question are shown in Figure 8. Note that BlockPy was referred to as "Blockly", and the use of the Spyder environment was referred to as "Python". These results suggest that students valued their experiences with BlockPy more than they did with NetLogo, but mostly felt that they were not getting enough Python experience. This is backed up by the qualitative data, where some students say "More Blockly, Less Python", but others ask for "More Blockly and More Python".

C. Usage of BlockPy

Over the six days spent using BlockPy, students were tasked with 40 classwork questions and 19 homework questions. Students ran their code an average of four times per problem (standard deviation is 1.8 times).

Students were asked if they felt successful in the transition from BlockPy to Spyder. 65% of the class agreed or strongly agreed, suggesting that there was still a sizeable population that still felt uncomfortable during that transition. The original design of the Mutual Language Translation featured the block and text view simultaneously, side-by-side. However, analysis of the logs reveals that most students did not take advantage of the feature. Only 5 students (roughly 15%) had used the conversion functionality at all, and fewer used it consistently. It is possible that students were observing the code as it changed, but they were not writing textual code. It is difficult to say why exactly students did not take advantage of it, and the qualitative data are not helpful. Our current hypothesis is that students were confused by the interface, which required manual conversion to go from text to blocks. In our new version, the conversion happens automatically, simply by switching tabs, and we will provide intentional opportunities for the students to switch.
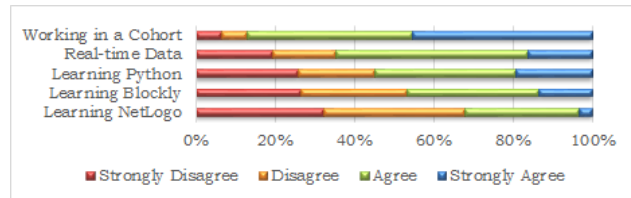
Students were surveyed about what helped their learning the most, and the results are shown in Figure 9. Peer learning and the instructors were about on par with the automatic feedback given in BlockPy, suggesting the strong value in using such a system. Despite the popular response to the Data Explorer, relatively few students took advantage of it (11 students, roughly 31%). Since more than 50% of the class reported finding value in the data explorer, it is possible that the students benefited from instructor presentations of the tool, even if they didn't take advantage of it themselves.

D. Data Science Context

Students were surveyed about their perceptions of the value of different course experiences with regards to their long-term career goals (Figure 10) and their interest in potential contexts for an introductory computing course (Figure 11). Both surveys suggest that students find data science to be compelling, but should be taken with a grain of salt, since students have negligible experience with alternative contexts, despite the care
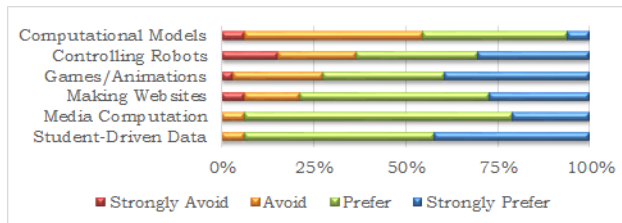
Figure 11. Students preference for different potential introductory contexts.

taken in wording the question. However, our preliminary results are exciting since they suggest that this is an approach worth exploring further.

IV. FUTURE WORK

BlockPy represents an ongoing research project and evolving practioners' tool. We have a number of features planned to expand the support for Python, in particular more blocks for data science explorations. We are also planning on expanding support for the guided feedback API for instructors and the support available through the compiler, such as leveraging static/dynamic type inference techniques to improve block rendering and error reporting.

We also have a number of research questions posed by the block-based nature of the interface. One of the biggest values of a block-based environment is that it can immediately expose the breadth of a rich API. This greatly reduces students' dependency on documentation. Of course, exposing this breadth can also be a downside, as students might be overwhelmed by the number of features in the interface. It is an open research question to decide what language features to expose and what rate to expose them at.

One of the major advantages of game and animation design as an introductory context is that they can make abstract concepts concrete. Further analysis is needed to determine the trade-offs of using different contexts. BlockPy can support this by supporting these alternative contexts, such as turtle graphics and media computation libraries.

Despite the substantial data collected in our pilot study, it is difficult to derive conclusive results due to the small population size and the evolving nature of BlockPy. At the time of submission for this publication, we are just finishing up a second pilot study with the latest version of the BlockPy interface; preliminary results suggest that recent improvements have overcome a number of limitations to the environment and user feedback has dramatically improved. Soon, we are conducting follow-up studies on the logged students' code, even as we collect more data on the newest iteration. We are hopeful that BlockPy will increase its user base, providing a larger sample of learners to conduct research on, and provide more meaningful data. We have concrete plans to host a larger study this Spring 2016.

A. Missing Language features

BlockPy is being developed in an on-demand fashion, driven by immediate user-needs. Unfortunately, this suggests a number of limitations. For example, the block interface does not support a number of advanced Python features, such as an interface for writing Object-oriented Classes. A number of other features are omitted too, at the time of writing: tuples, list comprehensions, and while loops, for example. This does not mean that students cannot write programs featuring classes or these features. Python code using these features will render in BlockPy as embedded text blocks, rather than as regular blocks, and will execute using Skulpt perfectly normally. There is no technical impediment to supporting these features, the process is limited only by time and community interest.

B. Missing Contexts

Environments aimed at younger learners often revolve around the creation of games and animations. Motivationally speaking, they focus on fun and situationally interesting experiences rather than practical, useful experiences. Although our primary vision is for users to learn in the context of data science projects, we believe that there is room for computational creativity. The underlying Python execution environment, Skulpt, has built-in support for turtle graphics (where an agent is directed around a drawing canvas), the Processing graphics library, and even HTML DOM manipulation. Although we have not exposed these APIs through the Block interface yet, we have a clear plan to do so. Ideally, these tools can be used to supplement lessons using data sets; students can create computational art based real-world data, marrying the art and science. Consider, for example, a program that draws a different scene based on the weather outside. Simple improvements to the "Sharing" system could allow students to interact with the client HTML DOM to create a stand-alone web application where users can choose a location and see the weather there.

V. CONCLUSION

In this paper, we have introduced our block-based environment for Python, named BlockPy. It is open-source and available for use for free at http://www.blockpy.com/. We believe that BlockPy represents a new paradigm for introductory learners, blending interactive support with a strong path to programming maturity. By teaching in the context of data science, we can

provide authenticity even as we move students out of the system towards a more serious environment. Research with this environment will help answer crucial questions about the value of data science and blocks. Our hope is that BlockPy's open-nature can encourage learners from diverse fields to engage with computing in a way that will lead to a computing-rich future for a larger population.

REFERENCES

[1] A. Vihavainen, J. Airaksinen, and C. Watson, "A systematic review of approaches for teaching introductory programming and their influence on success," in *Proceedings of the Tenth Annual Conference on International Computing Education Research*, ser. ICER '14. New York, NY, USA: ACM, 2014, pp. 19–26. [Online]. Available: http://doi.acm.org/10.1145/2632320.2632349

[2] B. D. Jones, "Motivating students to engage in learning: The MUSIC model of academic motivation," *International Journal of Teaching and Learning in Higher Education*, vol. 21, no. 2, pp. 272–285, 2009.

[3] M. Guzdial and A. E. Tew, "Imagineering inauthentic legitimate peripheral participation: an instructional design approach for motivating computing education," in *Proceedings of the second international workshop on Computing education research*. ACM, 2006, pp. 51–58.

[4] P. Guo, "Python is now the most popular introductory teaching language at top us universities," *BLOG@CACM, July*, 2014.

[5] A. Ko, "Programming languages are the least usable, but most powerful human-computer interfaces ever invented," http://blogs.uw.edu/ajko/2014/03/25/programming-languages", 2014.

[6] T. W. Price and T. Barnes, "Comparing textual and block interfaces in a novice programming environment," in *Proceedings of the Eleventh Annual International Conference on International Computing Education Research*, ser. ICER '15. New York, NY, USA: ACM, 2015, pp. 91–99. [Online]. Available: http://doi.acm.org/10.1145/2787622.2787712

[7] D. Weintrop and U. Wilensky, "Using commutative assessments to compare conceptual understanding in blocks-based and text-based programs," in *Proceedings of the Eleventh Annual International Conference on International Computing Education Research*, ser. ICER '15. New York, NY, USA: ACM, 2015, pp. 101–110. [Online]. Available: http://doi.acm.org/10.1145/2787622.2787721

[8] M. Guzdial, "Icer 2015 report: Blocks win - programming language design = = ui design," https://computinged.wordpress.com/2015/08/17/icer-2015-blocks-win-programming-language-design-ui-design/, 2015.

[9] T. Tang, S. Rixner, and J. Warren, "An environment for learning interactive programming," in *Proceedings of the 45th ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '14. New York, NY, USA: ACM, 2014, pp. 671–676. [Online]. Available: http://doi.acm.org/10.1145/2538862.2538908

[10] S. H. Edwards, D. S. Tilden, and A. Allevato, "Pythy: Improving the introductory python programming experience," in *Proceedings of the 45th ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '14. New York, NY, USA: ACM, 2014, pp. 641–646. [Online]. Available: http://doi.acm.org/10.1145/2538862.2538977

[11] P. J. Guo, "Online python tutor: Embeddable web-based program visualization for cs education," in *Proceeding of the 44th ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '13. New York, NY, USA: ACM, 2013, pp. 579–584. [Online]. Available: http://doi.acm.org/10.1145/2445196.2445368

[12] D. Bau, M. Dawson, and A. Bau, "Using pencil code to bridge the gap between visual and text-based coding (abstract only)," in *Proceedings of the 46th ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '15. New York, NY, USA: ACM, 2015, pp. 706–706. [Online]. Available: http://doi.acm.org/10.1145/2676723.2678293

[13] Y. Matsuzawa, T. Ohata, M. Sugiura, and S. Sakai, "Language migration in non-cs introductory programming through mutual language translation environment," in *Proceedings of the 46th ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '15. New York, NY, USA: ACM, 2015, pp. 185–190. [Online]. Available: http://doi.acm.org/10.1145/2676723.2677230

[14] M. Felleisen, R. B. Findler, M. Flatt, and S. Krishnamurthi, "How to design programs," 2001.

[15] A. C. Bart, "Situating computational thinking with big data: Pedagogy and technology (abstract only)," in *Proceedings of the 46th ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '15. New York, NY, USA: ACM, 2015, pp. 719–719. [Online]. Available: http://doi.acm.org/10.1145/2676723.2693616

[16] D. Parsons and P. Haden, "Parson's programming puzzles: a fun and effective learning tool for first programming courses," in *Proceedings of the 8th Australasian Conference on Computing Education-Volume 52*. Australian Computer Society, Inc., 2006, pp. 157–163.