

Self-Paced, Instructor-Assisted Approach to Teaching Linear Algebra

Pavel Solin

Abstract. In this paper we explain why we decided to abandon the traditional classroom instruction of Linear Algebra and switched to a self-paced, instructor-assisted course format. We also describe our experience with creating a self-paced course in NCLab, and with teaching it at a tier-1 research university for three consecutive semesters. The results of a student survey and student comments are presented.

Mathematics Subject Classification (2010). Primary 97D40; Secondary 97H60, 97U50.

Keywords. Linear algebra, mathematics, higher education, self-paced course, learning by doing.

1. Introduction

After more than 10 years of teaching Linear Algebra in the traditional way and seeing students struggle with the pace of the course every single semester, we decided to develop a hands-on self-paced course, stop lecturing, and let the students learn by doing. This decision was based on our long-term successful experience with using self-paced, instructor-assisted courses in other areas including computer programming, 3D modeling, and data analytics.

The new course is public¹. Any instructor can create a copy and use it as-is, or adjust it for his or her needs. The course allows students to learn actively, by reading and solving problems rather than by sitting, listening, and taking notes. Students work through detailed interactive online coursework, and they must demonstrate mastery of each concept before advancing to the next one. Their understanding of the material is validated each step of the way by instant feedback from the software platform. In the classroom, the instructor interacts with students one-to-one and helps them learn most efficiently. In remote instruction, NCLab² provides the instructor with a detailed real-time

¹Request free access to the course at <http://nclab.com/courses/linear-algebra/>.

²NCLab is an educational cloud platform available at <http://nclab.com/>.

overview of each student's progress, so again s/he can assist each student most efficiently.

We have been using the course both in the classroom and for online instruction with remarkable success. The majority of the students became more engaged, they learned the material better, enjoyed the course more, and the average scores have improved by 21% compared to our prior traditional classroom instruction. In this paper we share our experience with creating the coursework, as well as with using it as a self-paced, instructor-assisted course at a tier-1 research university.

1.1. Linear Algebra

Linear Algebra is one of the fundamental pillars of mathematics and computer science. It studies systems of linear equations and the related concepts of vectors, matrices, linear transformations and linear spaces. It has direct applications in science, engineering, data analytics, machine learning, artificial intelligence, business, finance, and many other areas. It is also indispensable in other areas of mathematics including calculus, ordinary and partial differential equations, numerical analysis, graph theory, etc.

Amazingly, many real-life objects such as images, 3D models, or large datasets can be represented by matrices. Then, techniques of Linear Algebra can be used to compress images, recognize people's faces in photographs, manipulate 3D objects on a computer screen, clean large datasets and extract important information from them, etc. Last but not least, studying Linear Algebra develops our thinking in ways that other branches of mathematics do not. Therefore, it is important that all students of mathematics, science, engineering, and related areas study Linear Algebra.

1.2. Limitations of Traditional Classroom Instruction

For clarity, by "traditional classroom instruction" we mean a lectured class where one instructor talks and N students sit, listen, take notes, and occasionally ask questions. The limitations of traditional classroom instruction have been discussed directly or indirectly in many studies related to the flipped classroom model. An extensive amount of references can be found in the nice survey article [3]. In the following, we share our personal experience with traditional classroom instruction which ultimately lead us to abandoning it, and replacing it with a self-paced, instructor-assisted course.

In our view, the most serious limitation of the traditional classroom setting is that students are not able to learn at their own pace. The learning styles and abilities of students vary. The instructor tries to cope with that by trying to find a pace which keeps the most students engaged. However, even if an optimal pace is achieved, it cannot work for all students. Some students will find it too fast, and the instructor will quickly lose them. In our experience, this is a significant amount. Such students will end up taking notes without understanding, and then tediously relearning the material later on their own. Or, they will just use Chegg [2] and purchase access to answers

to homework problems. As a side note, Cheating-as-a-Service (CaaS) has become a real problem of today's higher education. According to MacroTrends [6], Chegg is a \$10 billion enterprise.

In the traditional classroom, students are exposed to several new concepts per lecture. They are unable to practice and master one concept at a time before moving on to the next one. It should be noted that students do not learn by listening or taking notes. To learn a new concept, they need to practice it and validate their understanding by receiving positive feedback. Or, by being advised that they made a mistake, being given a chance to correct their mistake, and then receiving positive feedback. The importance of instant feedback in the learning process can't be stressed enough, yet it is not part of traditional classroom instruction.

Last but not least, traditional classroom instruction teaches students to be passive. Instead of being rewarded for taking the initiative and learning actively, students get used to the fact that it is normal to sit and wait for somebody to take care of them and teach them. They develop neither independent problems solving skills nor good communication skills. However, when they leave school and enter the workplace, they will find themselves in a shockingly different environment. They will be given tasks they don't know how to solve, and they will be expected to solve them, learning on their own whatever they don't know. Indeed.com [5] lists the lack of independent problem solving skills and good communication skills as the main reasons why college graduates struggle at their first job.

Switching to a self-paced, instructor-assisted course changed the atmosphere in our classroom completely. The students accepted the responsibility for their success. They became more active, communicated with the instructor and with their peers, sometimes they even worked together. Ultimately they became more successful. We have used the self-paced, instructor assisted course format for three consecutive semesters so far, and have experienced the same positive outcomes each time. Also the instructor gets rewarded, by seeing more students succeed. Finally, our student evaluations improved dramatically compared to the previous years of traditional classroom teaching.

1.3. The NCLab Matrix App

Before we share our experience with creating and using a self-paced Linear Algebra course [9], let's describe one particular widget which stands out by its practical usefulness: the NCLab Matrix App [13].

One of the least popular aspects of Linear Algebra is the manipulation of matrices via elementary row operations. This is a tedious and time-consuming process where students frequently make numerical mistakes. Yet, elementary row operations play a crucial role in Linear Algebra. They are needed to solve linear systems via the Gauss elimination and to invert matrices. They are also needed to transform matrices into their echelon and reduced echelon forms. These forms in turn reveal important information about the linear system including pivot and non-pivot columns and rows, basic and free variables, existence and uniqueness of solution, injectivity and surjectivity of the linear

map associated with the matrix, dimension of the null and column spaces, etc. Unfortunately, the cumbersome nature of performing elementary row operations with pencil and paper prevents many students from becoming truly familiar with these essential concepts.

For illustration, let's consider a sample four-equation system,

$$\begin{aligned} -3x_2 - 6x_3 + 4x_4 &= 9 \\ -x_1 - 2x_2 - x_3 + 3x_4 &= 1 \\ -2x_1 - 3x_2 + 3x_4 &= -1 \\ x_1 + 4x_2 + 5x_3 - 9x_4 &= -7 \end{aligned}$$

which can be written in the matrix form $Ax = b$,

$$\begin{pmatrix} 0 & -3 & -6 & 4 \\ -1 & -2 & -1 & 3 \\ -2 & -3 & 0 & 3 \\ 1 & 4 & 5 & -9 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} = \begin{pmatrix} 9 \\ 1 \\ -1 \\ -7 \end{pmatrix}$$

The augmented matrix $(A|b)$ has the form

$$\left(\begin{array}{cccc|c} 0 & -3 & -6 & 4 & 9 \\ -1 & -2 & -1 & 3 & 1 \\ -2 & -3 & 0 & 3 & -1 \\ 1 & 4 & 5 & -9 & -7 \end{array} \right)$$

In order to obtain the corresponding reduced echelon form, one needs to perform 12 elementary row operations:

$$\left(\begin{array}{cccc|c} 0 & -3 & -6 & 4 & 9 \\ -1 & -2 & -1 & 3 & 1 \\ -2 & -3 & 0 & 3 & -1 \\ 1 & 4 & 5 & -9 & -7 \end{array} \right) \rightarrow \left(\begin{array}{cccc|c} 1 & 4 & 5 & -9 & -7 \\ -1 & -2 & -1 & 3 & 1 \\ -2 & -3 & 0 & 3 & -1 \\ 0 & -3 & -6 & 4 & 9 \end{array} \right) \rightarrow$$

$$\left(\begin{array}{cccc|c} 1 & 4 & 5 & -9 & -7 \\ 0 & 2 & 4 & -6 & -6 \\ -2 & -3 & 0 & 3 & -1 \\ 0 & -3 & -6 & 4 & 9 \end{array} \right) \rightarrow \left(\begin{array}{cccc|c} 1 & 4 & 5 & -9 & -7 \\ 0 & 2 & 4 & -6 & -6 \\ 0 & 5 & 10 & -15 & -15 \\ 0 & -3 & -6 & 4 & 9 \end{array} \right) \rightarrow$$

$$\left(\begin{array}{cccc|c} 1 & 4 & 5 & -9 & -7 \\ 0 & 1 & 2 & -3 & -3 \\ 0 & 5 & 10 & -15 & -15 \\ 0 & -3 & -6 & 4 & 9 \end{array} \right) \rightarrow \left(\begin{array}{cccc|c} 1 & 4 & 5 & -9 & -7 \\ 0 & 1 & 2 & -3 & -3 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & -3 & -6 & 4 & 9 \end{array} \right) \rightarrow$$

$$\left(\begin{array}{cccc|c} 1 & 4 & 5 & -9 & -7 \\ 0 & 1 & 2 & -3 & -3 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & -3 & -6 & 4 & 9 \end{array} \right) \rightarrow \left(\begin{array}{cccc|c} 1 & 4 & 5 & -9 & -7 \\ 0 & 1 & 2 & -3 & -3 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -5 & 0 \end{array} \right) \rightarrow$$

$$\left(\begin{array}{cccc|c} 1 & 4 & 5 & -9 & -7 \\ 0 & 1 & 2 & -3 & -3 \\ 0 & 0 & 0 & -5 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right) \rightarrow \left(\begin{array}{cccc|c} 1 & 4 & 5 & -9 & -7 \\ 0 & 1 & 2 & -3 & -3 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right) \rightarrow$$

$$\left(\begin{array}{cccc|c} 1 & 4 & 5 & -9 & -7 \\ 0 & 1 & 2 & 0 & -3 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right) \rightarrow \left(\begin{array}{cccc|c} 1 & 4 & 5 & 0 & -7 \\ 0 & 1 & 2 & 0 & -3 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right) \rightarrow \left(\begin{array}{cccc|c} 1 & 0 & -3 & 0 & 5 \\ 0 & 1 & 2 & 0 & -3 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right)$$

The reduced echelon form reveals that:

- x_1, x_2, x_4 are basic variables and x_3 is a free variable.
- The linear system is consistent and has infinitely many solutions.
- The linear map $T(u) = Au$ is not injective because there exist non-pivot columns.
- The linear map $T(u) = Au$ is not surjective because there exist non-pivot rows.
- The coefficient matrix A has a one-dimensional null space and three-dimensional column space, etc.

Finally, the reduced echelon form can be used to express the infinite solution set in parametric vector form,

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} = \begin{pmatrix} 5 \\ -3 \\ 0 \\ 0 \end{pmatrix} + x_3 \begin{pmatrix} 3 \\ -1 \\ 1 \\ 0 \end{pmatrix}$$

We are mentioning all these elementary facts because we identified the lack of practice as the main reason why many students have difficulties with the aforementioned concepts. Unfortunately, it is not possible to practice this analysis in the classroom enough, using just pencil and paper. Out of curiosity, we timed our students, and it took over 15 minutes until all of them were able to obtain this reduced echelon form. This included debugging the elementary row operations for some students who were unable to find their numerical mistakes quickly.

For this reason, NCLab provides an interactive web-based Matrix App which any course creator can use. The Matrix App allows students to easily define matrices on their devices, and perform elementary row operations with their fingers or with the mouse/keyboard. The Matrix App uses HTML5 [4] and Mathjax [7] on the frontend and Python [10] on the backend. It satisfies AA-level web accessibility requirements [1] and is fully operable using keyboard only. It works on desktop computers, laptops, chromebooks, and on touch devices. The Matrix App makes it possible to enter standard matrices, as well as augmented matrices with one or more right-hand sides. It also provides a Python shell where students can use Numpy, Scipy, and other

Python libraries. This is described in more detail in a separate document [13].

For illustration, let's use the Matrix App to solve the same task. It takes less than a minute to set the dimension of the matrix and enter the values using either fingers, mouse or keyboard:

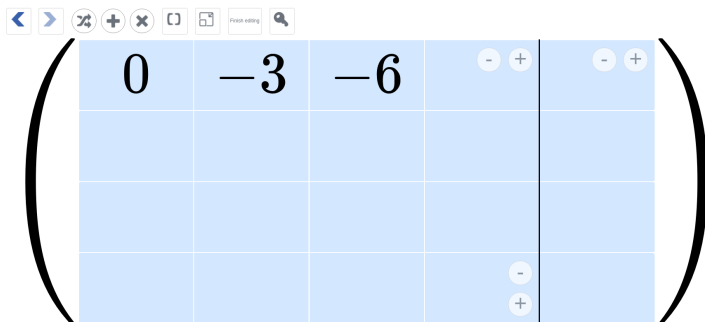


FIGURE 1. Creating the matrix

This is the finished augmented matrix $(A|b)$:




FIGURE 2. Augmented matrix

It takes less than another minute to perform the 12 elementary row operations, transforming the augmented matrix into its reduced echelon form shown in Fig. 3. There is zero chance to make a numerical mistake on the way. The YouTube video [12] shows the full procedure in real time.

$$\left(\begin{array}{cccc|c} 1 & 0 & -3 & 0 & 5 \\ 0 & 1 & 2 & 0 & -3 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{array} \right)$$

FIGURE 3. Reduced echelon form

The Matrix App uses a simple on-screen keypad (shown in Fig. 4) to enter matrix values and coefficients for scaling and replacement operations.

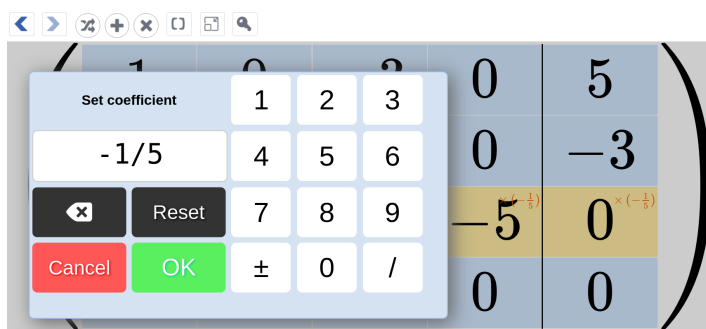


FIGURE 4. On-screen keypad used to enter matrix entries, as well as coefficients for scaling and replacement operations

In summary, the Matrix App makes it easy for students to manipulate matrices via elementary row operations, solve linear systems via the Gauss elimination, invert matrices, and create echelon and reduced echelon forms. This makes it easier for the students to learn essential Linear Algebra concepts related to the solvability of linear systems, properties of the corresponding linear maps, properties of null and column spaces, structure of infinite solution sets, etc.

2. Creating a Self-Paced Course

We decided to develop the self-paced Linear Algebra course using the NCLab [8] platform. Any instructor is welcome to make a copy of the course and either use it as-is, or tweak it for his or her own purposes. NCLab is an educational cloud platform which allows educators to create high-quality self-paced online coursework with emphasis on practice-based hands-on learning. Previous courses developed with NCLab include 3D modeling, coding, data

analytics, geometry, and math³. New courses are created by combining text, \LaTeX formulas, images, videos, examples, practice exercises with solutions, and autograded practical tasks and quizzes. Let's have a closer look!

2.1. Tutorials

Tutorials are created using the TinyMCE text editor [14] which supports Mathjax for entering \LaTeX formulas. The TinyMCE editor makes it easy to include images, videos and HTML links. It also provides an ADA web accessibility checker which is easy to use and ensures that the course content is web accessible.

When creating course contents, we prefer to split tutorials into short self-contained pages, rather than creating long pages which the students can scroll through quickly. Students then use arrows to move between pages. Sometimes they have to solve a problem or answer a question before they are allowed to move to the next page.

Our course has 25 Sections and each Section has 10 Levels. For illustration, let's go through Section 4 which introduces elementary row operations, their reversibility, augmented matrix, Gauss elimination, and row-equivalent matrices. Level 4.1 begins with an outline of Section 4:

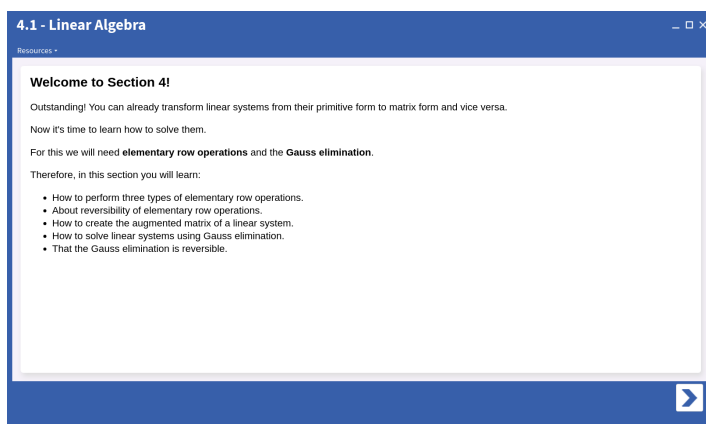


FIGURE 5. Outline of Section 4.

The next page informs the students that there are three types of elementary row operations:

³An overview of these courses can be found at <http://nclab.com/courses/>.

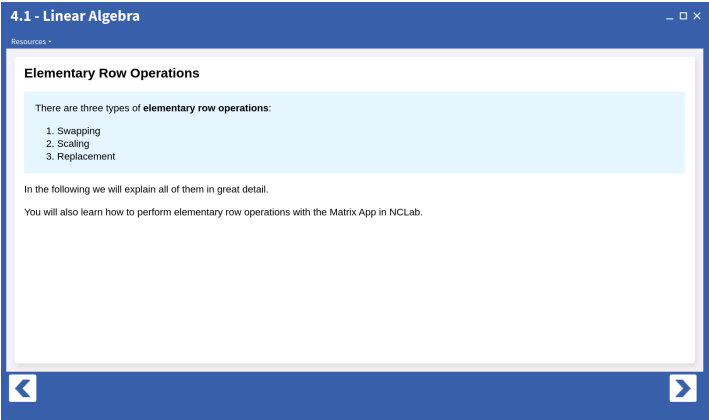


FIGURE 6. Three types of elementary row operations.

The third page of Level 4.1 introduces the swapping row operation:

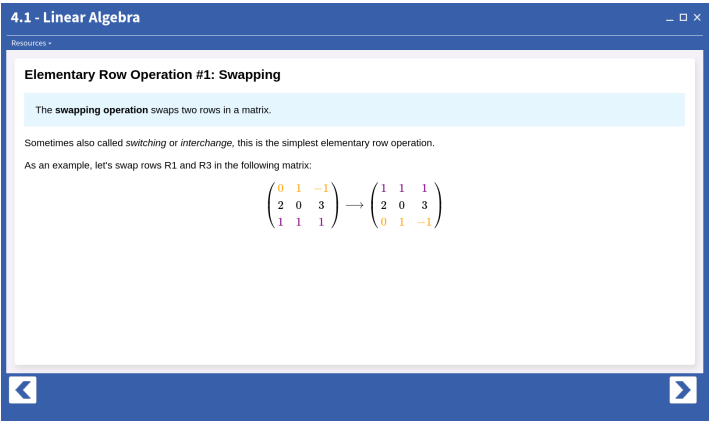


FIGURE 7. The swapping row operation.

Fig. 8 shows the same page opened in editing mode:

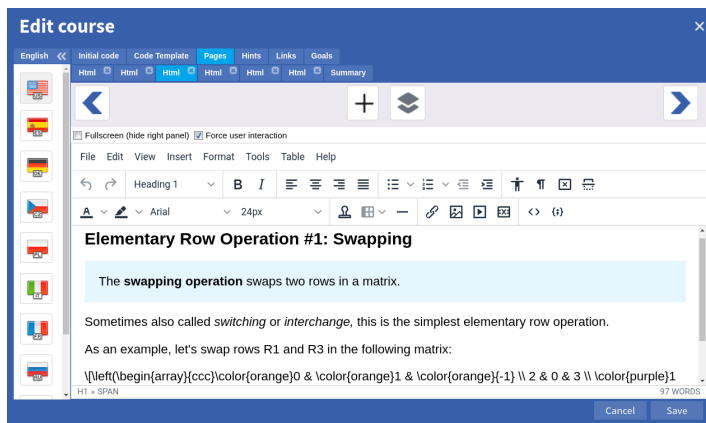


FIGURE 8. Third tutorial page of Level 4.1 in editing mode.

The fourth page of Level 4.1 shows reversibility of the swapping row operation:

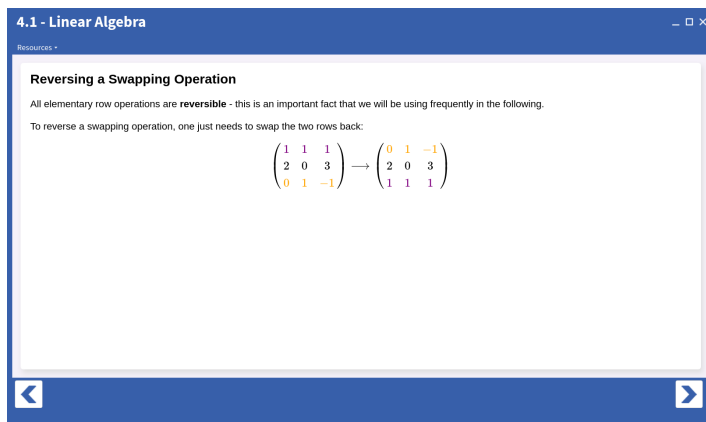


FIGURE 9. Reversibility of the swapping row operation.

The fifth page features a short video which explains how to perform the swapping row operations in the Matrix App:

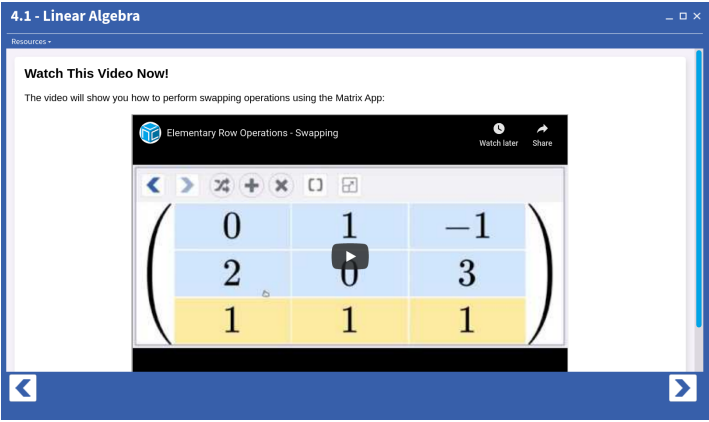


FIGURE 10. Swapping rows in the Matrix App.

The sixth page formulates the practical task for this Level:

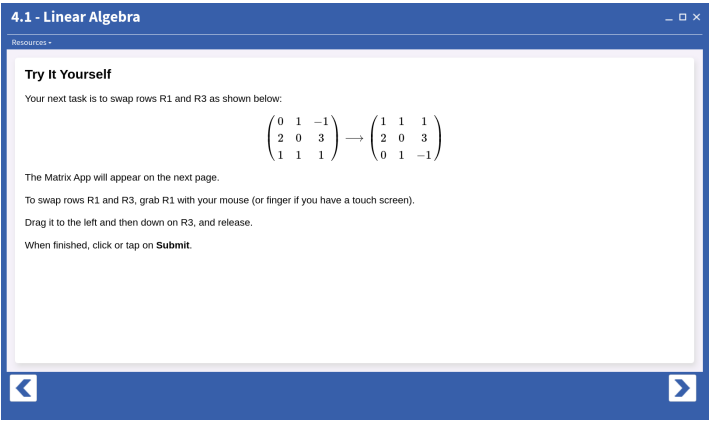


FIGURE 11. Practical task for this Level.

The last page of Level 4.1 will open the Matrix App:




FIGURE 12. The matrix opens in the Matrix App.

If the student fails to perform the swapping operation correctly, s/he is not allowed to proceed to the next Level. A helpful feedback is provided so that the student can see where s/he made a mistake:




FIGURE 13. The student failed to perform the swapping operation correctly.

Only after solving the task correctly, the student is allowed to proceed to Level 4.2:

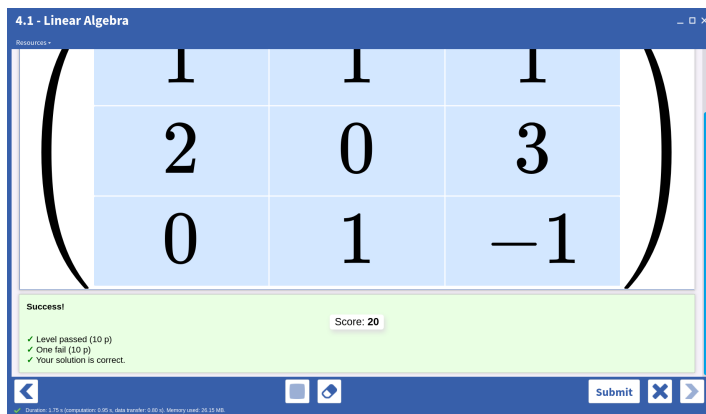


FIGURE 14. Student solved the task correctly and can proceed to the next Level.

Here is an overview of what the student learns in the rest of Section 4:

- Level 4.2: Scaling operation + practical task.
- Level 4.3: Replacement operation + practical task.
- Level 4.4: Augmented matrix, Gauss elimination, row-equivalent matrices + practical task.
- Level 4.5: Gauss elimination for a 2×2 matrix + practical task.
- Level 4.6: Reversibility of Gauss elimination + practical task.
- Level 4.7: Gauss elimination for a 3×3 matrix + practical task.
- Level 4.8: Practical task: Gauss elimination for a 3×3 matrix.
- Level 4.9: Practical task: Gauss elimination for a 4×4 matrix.
- Level 4.10: Review exercises + quiz.
- Level 4.11: Review guide for Section 4.

As the reader can see, anything new that the student learns is followed by an autograded practical task where they can validate their understanding of the new concept. We strongly believe that it is essential for the students to practice each new concept, and have their understanding validated, before they advance to the next one. As we mentioned back in Section 2, this is something that traditional classroom instruction does not allow, and on the other hand, that self-paced courses are perfect for.

2.2. Review Exercises and Quizzes

Section 4 of the course which was presented above is exceptional in that it only contains one Quiz. Most other Sections contain five tutorial Levels and each of them is followed by Review Exercises and a Quiz. For instance, in Section 5 of the course students learn about the echelon and reduced echelon forms, and how to express infinite solution sets in parametric vector form. The following Fig. 15 shows page 1 of 3 of Review Exercises preparing the students for Quiz 5.10:

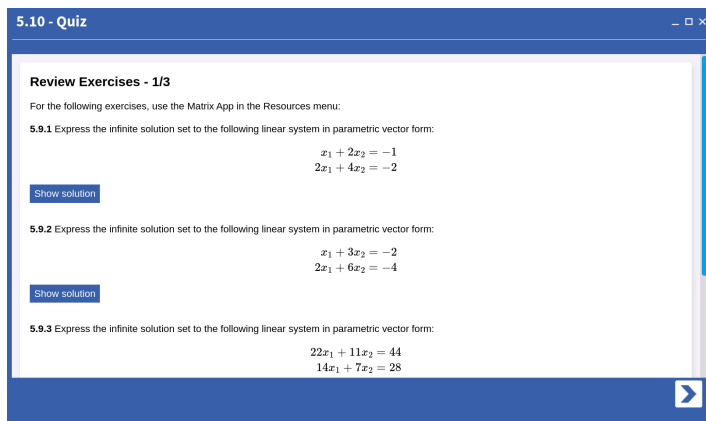


FIGURE 15. Review Exercises with solutions, preparing the students for Quiz 5.10.

Clicking or tapping on the Solution button reveals a detailed step-by-step explanation of the solution procedure:

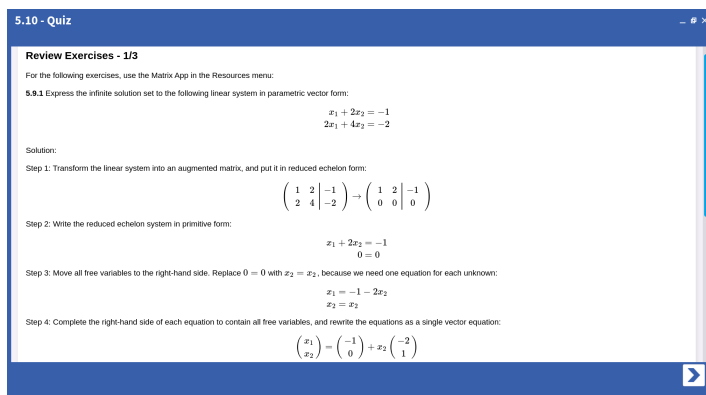


FIGURE 16. Review Exercises are accompanied with detailed explanations of the solution procedure.

Now let's talk about Quizzes. NCLab allows course creators to design auto-graded Quizzes with questions that expect one of the following 11 types of answers:

1. True or False.
2. Number.
3. Open Response.
4. Multiple Bchoice.
5. Checkboxes.
6. Computer Code.

7. Vector or Matrix
8. Eigenvectors.
9. Basis.
10. Orthogonal Basis.
11. Ordering.

While Open Response, Computer Code and Ordering are not used in our Linear Algebra course, all other types of questions are used. The first seven types of questions are self-explanatory. Question types 8 - 11 do not have a unique answer. For instance, Eigenvectors are determined up to a multiplicative constant, they can be ordered differently, and eigenvectors corresponding to the same eigenvalue may come as different linear combinations. The autograding algorithm accounts for all this.

In our Linear Algebra course, each Quiz consists of five questions. Sometimes, a significant amount of work is required to answer a question. The following Fig. 17 shows the first part of Quiz 5.10 which is dedicated to echelon and reduced echelon forms, and expressing infinite solution sets in parametric vector form.

The screenshot shows a quiz interface titled "5.10 - Quiz". It contains two questions, each worth 20 points.

Question 1 (20 points)
 Transform the 2×4 matrix to its reduced echelon form!

$$\begin{pmatrix} -2 & 3 & -2 & 3 \\ 3 & 2 & 3 & 2 \end{pmatrix}$$

(Link to Matrix App)

Below the question is a row of input fields with "+" and "-" icons for adjusting dimensions.

Question 2 (20 points)
 Transform the 3×4 matrix to its reduced echelon form!

$$\begin{pmatrix} 1 & 2 & 3 & -1 \\ 4 & 5 & 6 & 3 \\ 7 & 8 & 9 & 5 \end{pmatrix}$$

(Link to Matrix App)

At the bottom right, there is a "Submit" button and navigation arrows.

FIGURE 17. Quiz 5.10.

These questions require an answer of type 7 (Vector or Matrix). To answer them, the student first uses the + and - icons to set the dimensions of the vector or matrix, and then s/he has to enter the correct values.

Fig. 18 shows the same Quiz in editing mode, with the master solution entered by the course creator. When the student submits the Quiz, his or her solution is compared with the master solution:

It is possible to provide several versions of each question, of which one is selected randomly when the student opens the Quiz. Also, Quizzes can be retaken after 12 hours. When a student retakes a Quiz, s/he sees different questions (if at least two versions of the question are provided by the course creator).

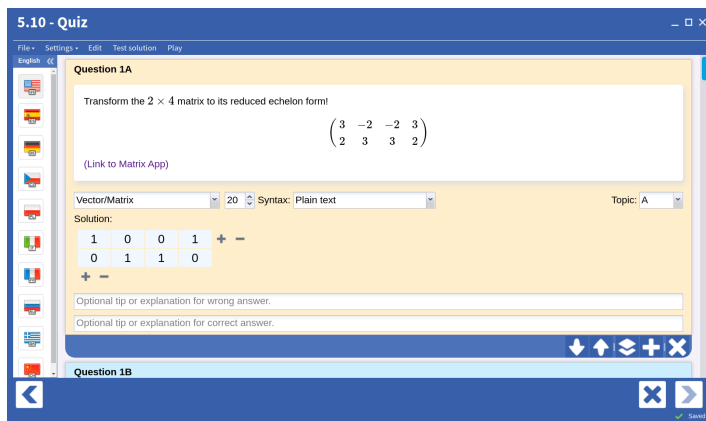


FIGURE 18. Quiz 5.10 in editing mode.

2.3. Integrating Computation Into the Course

NCLab's backend provides Python, Java, Javascript and other programming languages, GNU Octave, Scilab, \LaTeX , PLaSM, SCAD, and other software related to scientific computing and visualization. This makes it easy for the course creator to incorporate computing into his or her course. We use Python throughout our entire Linear Algebra course. With every concept taught, starting with adding and subtracting vectors, multiplying them with scalars, and performing the dot product, we always show the students how this is done with Numpy or Scipy:




FIGURE 19. Performing the dot product of vectors with Numpy.

We also use slightly more advanced Python programs to illustrate selected concepts. For instance, when the students learn about rotation matrices in Section 8 of the course, they can experiment with a simple program which rotates 2D objects defined via vertex vectors:

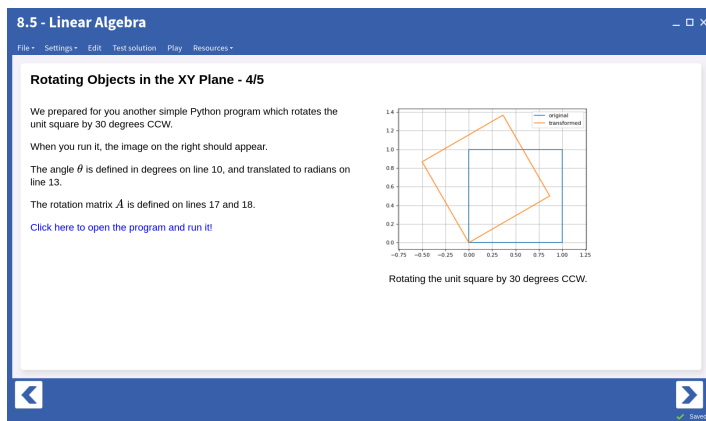


FIGURE 20. Students are asked to run a Python program which rotates objects in the XY plane.

The program defines a unit square using its vertex vectors, rotates them, and displays the transformed object. The students can then define their own 2D objects and rotate them:




FIGURE 21. Experimenting with planar rotations.

As another example, when the students learn about SVD in Section 23 of the course, they can experiment with a simple program which compresses images:

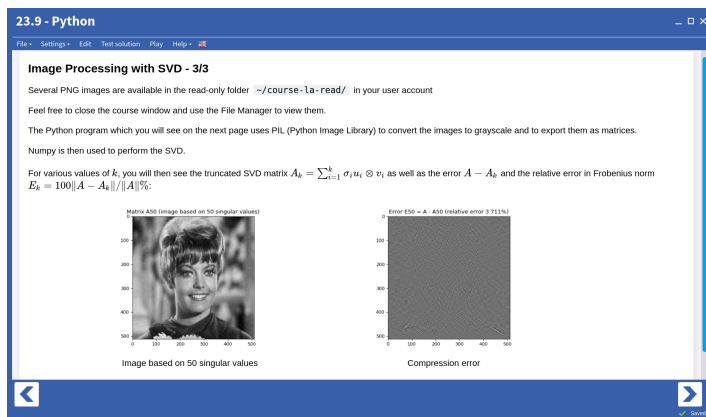


FIGURE 22. Using SVD to compress images.

In our course, the computational aspect is merely informative. Students are not required to use Python nor are they tested on it. Only the (optional) last Section 25 of the course is dedicated to Linear Algebra with Python. There, students learn how to use Numpy and Scipy to solve various tasks ranging from solving linear systems and eigenproblems to solving least squares problems, performing spectral decomposition, QR factorization, and SVD.

3. Teaching With the Self-Paced Course

In the following we will share our experience with teaching the Linear Algebra course. Before we dive into the details though, let's explain the difference between a self-paced, instructor-assisted course and flipped classroom.

3.1. Self-Paced, Instructor-Assisted Learning vs. Flipped Classroom

By definition, in a flipped classroom students learn at home on their own, and then in the classroom they practice what they learned [11]. This, however, does not solve the problem with lack of timely feedback as explained in the previous section. When students "learn" at home and have to wait several hours or even days until their understanding of the material is validated in the classroom, this is equally inefficient as the standard classroom instruction.

In contrast to this, a self-paced, instructor-assisted course is designed for the students to learn in the classroom, at their own pace and with the instructor's individual real-time assistance, practicing each new concept immediately, and receiving intelligent real-time feedback from a software platform. The immediate application of each concept learned, combined with instant feedback, is a key element of this instructional model. A self-paced, instructor-assisted course can be easily adapted to an online setting where students learn on their own, and have real-time access to an instructor via an online meeting platform.

3.2. Course Overview

The new Linear Algebra course is an equivalent of a full semester, three-credit course. The coursework consists of 25 Sections, each Section being equivalent to one lecture including the corresponding homework. One Section includes on average five tutorial levels, each followed by a practice level where students instantly check if they grasped the concept. The practice levels provide intelligent real-time feedback and help which motivates students to keep learning. Each practice level can be retaken one time after 12 hours. Students need approximately 3-4 hours to complete each Section, but the amount of time depends on their aptitude and learning habits.

3.3. Benefits for the Students

An instructor who is not lecturing is able to work with each student individually, helping them as much as needed. In our experience, a large part of the students do not require any help, but the instructor should check with them at least once per class anyway. Other students will form pairs or groups, communicate, and learn together. In summary:

- Students are not "getting lost" like in the traditional classroom.
- Students become more active, working through problems at their own pace with instructor's individual help.
- Instructor helps weaker students to perform better, and s/he can challenge the gifted ones to achieve more.
- Students can form pairs or groups and learn collaboratively.
- Classroom time is used more effectively: Students have to spend less time at home re-learning material and doing homework.

3.4. Benefits for the Instructor

Contrary to traditional lecture-based Linear Algebra courses, the instructor does not serve as the delivery vehicle for contents which students are perfectly capable of learning on their own. Instead, the instructor has more time to help every student individually as needed. The instructor has access to a detailed progress report of every student in real time:




FIGURE 23. Instructor sees each student's progress in real time.

The record shown in Fig. 23 corresponds to one student. The records are updated in real-time as the students progress through the course. Each line represents five Sections. The green checkmarks indicate completed tutorial levels, and the blue ones completed practice levels and quizzes. By clicking on the individual icons, the instructor obtains more details including the time spent, number of attempts needed, numerical score, and s/he can see precisely what mistakes the student made. In summary:

- The instructor does not deliver the material, but s/he helps the students on an individual basis.
- The instructor becomes more productive and achieves better results.
- The instructor has access to real-time student progress data.
- The instructor has detailed knowledge of how every student in class is doing.
- Students enjoy the class more than traditional lecture-based course.
- Classroom becomes a more active and productive environment.

3.5. Results of a Survey

Below we show a questionnaire which we give to the students every semester. So far we did it three times, and the outcomes were similar. The results shown below are from December 2019, at the end of the first semester we taught the course in the self-paced format:

What is your overall level of satisfaction with the course?

5 = very satisfied, 4 = more than satisfied, 3 = satisfied, 2 = partly, 1 = not at all

5 IIIIIIIII

4 IIIII

3 IIIIIII

2

1

How would you rate the difficulty of the material?

5 = very difficult, 4 = difficult, 3 = just right, 2 = easy, 1 = very easy

5 IIIII

4 IIIIIIIII

3 IIIIIII

2

1

Compare your classroom experience in this course to traditional math lecture courses. Does this course use your time more effectively?

5 = this hands-on course is much more productive, 4 = this hands-on course is somewhat more productive, 3 = the experiences are about equal, 2 = traditional course is somewhat more productive, 1 = traditional course is much more productive

5 IIIIIIIII

4 IIIIIII

3 IIIII

2 I

1

In other courses you take, would you like to learn using this hands-on method, or using the traditional lecture + homework format?

5 = strongly prefer the hands-on method used in this course, 4 = somewhat prefer the hands-on method used in this course, 3 = no preference, 2 = somewhat prefer the traditional lecture + homework, 1 = strongly prefer the traditional lecture + homework

5 IIIIIIIIIII

4 IIIII

3 IIII

2 I

1

3.6. Student Comments

Since we started using the self-paced course, students have been sending us emails expressing their satisfaction with the course format. Here are a few examples of their comments:

"Having never dealt with Linear Algebra, the format of this class has been extremely helpful in my success as a student. Being able to tackle the class at my own pace, collaborate with peers, and receive direct help from the professor in real time has been an invaluable experience." M. N.

"I love the self-taught aspect of the course supplemented by the teacher. Instruction is excellent. The flexibility is superb." A. R.

"This is how I like to learn, at my own pace. I like the bite-sized tutorials and I feel a real sense of accomplishment when I complete the exercises that follow them." M. J.

"I have really enjoyed the course so far, and based on my test score this is the most successful I have been in a math course since high school. I would not call other classes I have had difficult but simply unengaging, and this class has turned that around." Y. M.

"I cannot compliment enough any training that involves students actually working on the problems rather than just listening." A. S.

"The format of your Linear Algebra (MATH330) class using NCLab lets me have an easier time learning the material than in other courses. Thank you for having the class set up using this tool." M. S.

"The problem I have in most classes is that I tend to fall asleep during lectures, so being able to learn at my own pace while I'm most awake is a blessing." F. T.

"I really enjoyed the online/modular format of this course. The material was presented clearly and concisely and allowed me to focus longer on concepts that were difficult, while glancing over concepts I was already familiar with." A. C.

"The format of the course ensures that I am learning what I need to before moving onto the next section. I can learn at my own pace and find supplemented instruction as needed. In many of my previous math courses I have had to spend time at home relearning the material. I really enjoy this format." M. S.

"The most beneficial parts of this hands on course is the ability to access the instructor and the strait-to-the-point lecture and side notes." E. M.

"The feature I find most beneficial about this course format is that the notes are always one click away so I don't need to keep track of writing down my own notes. I also like how we are able to learn at our own pace." A. L.

Appendix A: Tips, Tricks, and Lessons Learned

Units 1 - 4 (Sections 1 - 20) correspond to a standard first Linear Algebra course. However, many students will be able to cover more ground based on their ability. Sections 21 - 25 discuss the Spectral Theorem, Least-Squares problems, QR factorization, SVD, large matrices (conditioning, memory issues), and linear algebra with Numpy and Scipy.

Students appreciate it when the instructor uses the first 10-15 minutes at the beginning of each class to explain the importance of what they will be learning today. After that, they can start working independently.

The course is designed to have an extremely smooth learning curve, and most students will go through it without requiring your assistance. As a consequence, class attendance will tend to drop. We experimented with various incentives aimed at increasing attendance, including extra credit for class attendance. But so far the method which seems to work best, is to make class attendance mandatory, unless the student completes the scheduled work in advance. We apply penalties if students break this rule.

To prevent students from falling behind, they are required to complete by the end of each calendar week all material which was scheduled for that week. We schedule one Section per class and the deadlines are in the syllabus. Penalties apply if students break this rule.

Some students are not very communicative, but we found it very important to have at least one conversation with each student per class. The same holds about exceptional students who seem to not require any attention. Without exception, every student appreciates your attention. The easy conversation starters are "Did you have any problems with the previous Section?", "Did you make any mistakes in the last quiz? Show me.". Discuss the difficulties and/or mistakes with them.

It is important to be absolutely supportive and avoid any judgment, even if it turns out that the student is struggling with some elementary concept from precalculus or even from high school. A number of times it happened to us that students who were not doing well gained confidence and improved dramatically after we helped them "sealed holes" in their knowledge of elementary concepts.

Students appreciate it if before each Midterm and Final they are given a detailed list of Levels to review. Some students will tend to race through the course and get far ahead. In our experience it is better to slow them down. This can be done by unlocking Sections gradually, or by informing the

students clearly what the plan for the week is, and tell them to not go beyond that.

We find that 80-90 % of the students are honest and want to do the required work and learn. There are no problems with these. But then there are a few who will try to avoid doing the work. We have been monitoring their behavior and improving the platform accordingly. This is an ongoing process. Last but not least, some students will lag and eventually drop out. The instructor dashboard will show you the students who have problems, and you can try to help them. But not all students want to be helped, and some of them will become unresponsive and quit no matter what you do.

Let the students fill out an anonymous questionnaire after they finish Unit 1. This may help you improve your teaching.

Regarding grading: The online coursework is autograded and you will see in percent how each student is doing. We usually count this as 40% of the final score. In addition to that, we usually do two midterms and one final exam per semester, each contributing with 20% to the final score. Exams are in person, in class. The NCLab platform does not provide online exams, but it provides thousands of Review Exercises with solutions, from which the instructor can randomly choose and build the exams. All Review Exercises are written in \LaTeX , therefore it is easy to extract the \LaTeX source of the exercises.

Last but not least, the course requires Internet access. Originally we started teaching it in a regular classroom (not in a computer lab). We were worried if every student will be able to use a laptop or tablet, but this was never a problem. Eventually, some students suggested that it would be great to move the class to a computer lab, which we did. But then we found that numerous students brought their own mobile devices anyway. Hence, it is not critically important to reserve a computer lab for this course. It is more important to be able to walk freely among the students. Therefore, large classroom with cinema-style seating should be avoided.

Appendix B: Course Syllabus

The course was designed to have extremely mellow learning curve, because in our experience, the students do not mind a slow pace. If some topic is discussed too slowly, then they can move through it quickly. But on the other hand, if a topic was explained insufficiently, then many students might struggle with it. In this sense, we prefer to err on the side of caution.

Compared to other Linear Algebra courses, our course may be discussing some aspects of Linear Algebra in more detail. For instance, concepts such as linear space, subspace, norm, inner product, orthogonal projection, etc. are explained in the context of vector spaces, polynomial spaces, and spaces of matrices. We also cover complex-valued linear spaces. It is the freedom of every instructor, however, to leave out the material they don't want to

include. Here is the course syllabus:

Section 1 (Vectors and Matrices I)

- What is an n -vector, how to add and subtract vectors, and multiply them with scalars (real numbers).
- What is the dot product of vectors.
- How is the dot product related to angle, orthogonality, Euclidean norm (length) and distance of vectors.
- What is an $m \times n$ (" m by n ") matrix and how to correctly determine matrix dimensions.
- How to add and subtract matrices, and multiply them with scalars.
- How to use indices when accessing individual entries in vectors and matrices.
- How to calculate matrix-vector and matrix-matrix products, and when they are defined and undefined.

Section 2 (Vectors and Matrices II)

- More about vectors and matrices, including zero vectors, zero matrices, identity matrices, and diagonal matrices.
- About commutative, associative, and distributive properties of vector and matrix operations.
- That some important vector and matrix operations are not commutative.
- How to count arithmetic operations involved in various vector and matrix operations.
- Which vector and matrix operations are most computationally demanding.
- About matrix transposition, and how to work with transposed and symmetric matrices.

Section 3 (Linear Systems)

- Basic facts about linear systems and their solutions.
- Basic terminology including equivalent, consistent and inconsistent systems.
- About the existence and uniqueness of solutions to linear systems.
- How to transform a linear system to a matrix form.

Section 4 (Elementary Row Operations)

- How to perform three types of elementary row operations.
- About reversibility of elementary row operations.
- How to create the augmented matrix of a linear system.
- How to solve linear systems using Gauss elimination.
- That the Gauss elimination is reversible.

Section 5 (Echelon Forms)

- What is an echelon form of a matrix and how to obtain it.
- How to use the echelon form to infer existence and uniqueness of solution.

- What is an under- and over-determined system.
- How to determine pivot positions and pivot columns.
- What are basic and free variables.
- What is the reduced echelon form and how to obtain it.
- About the uniqueness of the reduced echelon form.
- How to express solution sets to linear systems in parametric vector form.

Section 6 (Linear Combinations)

- What is a linear combination of vectors.
- How to write a linear system as a vector equation and vice versa.
- How to write a vector equation as a matrix equation $Ax = b$ and vice versa.
- What is a linear span, and how to check if a vector is in the linear span of other vectors.
- The Solvability Theorem for linear systems.
- That the linear system $Ax = b$ has a solution for every right-hand side if and only if the matrix A has a pivot in every row.

Section 7 (Linear Independence)

- What it means for vectors to be linearly independent or dependent.
- How to determine linear independence or dependence of vectors.
- How is linear independence related to linear span.
- How is linear independence related to the matrix equation $Ax = 0$.
- About singular and nonsingular matrices.
- About homogeneous and nonhomogeneous linear systems.

Section 8 (Linear Transformations I)

- Review of essential concepts including set, subset, relation, function, domain, codomain, image, and range.
- Review of basic properties of functions such as one-to-one (injective) functions, onto (surjective) functions, bijective functions, and inverse functions.
- Review of linearity and linear functions.
- About linear vector transformations of the form $T(x) = Ax$.
- How to figure out the matrix of a linear vector transformation.
- About rotations, scaling transformations, and shear transformations.
- That surjectivity of the transformation $T(x) = Ax$ is equivalent to the existence of solution to the linear system $Ax = b$ for any right-hand side.
- How to easily check if the transformation $T(x) = Ax$ is onto (surjective).
- That injectivity of the transformation $T(x) = Ax$ is equivalent to the uniqueness of solution to the linear system $Ax = b$.
- How to easily check if the transformation $T(x) = Ax$ is one-to-one (injective).
- That the solution to the linear system $Ax = b$ is unique if and only if the matrix A has a pivot in every column.

- That bijectivity of the transformation $T(x) = Ax$ is equivalent to the existence and uniqueness of solution to the linear system $Ax = b$ for any right-hand side.

Section 9 (Linear Transformations II)

- How to find a spanning set for the range of a linear transformation.
- What is the kernel of a linear transformation.
- How to find a spanning set for the kernel.
- That the kernel determines injectivity of linear transformations.
- About compositions of linear transformations.
- How to obtain the matrix of a composite transformation.
- About inverse matrices and inverse transformations.
- How to obtain the matrix of an inverse transformation.
- What can one do with inverse matrices.

Section 10 (Special Matrices and Decompositions)

- What are block matrices and how to work with them.
- How to count arithmetic operations in the Gauss elimination.
- How to work with block diagonal, tridiagonal and diagonal matrices.
- What are upper and lower triangular matrices.
- What are sparse and dense matrices.
- What are elementary matrices.
- Gauss elimination in terms of elementary matrices.
- Matrix inversion in terms of elementary matrices.
- About LU decomposition of nonsingular matrices.
- What is definiteness of symmetric matrices.
- How to check if a symmetric matrix is positive-definite, positive-semidefinite, negative-semidefinite, negative-definite or indefinite.
- About Cholesky decomposition of symmetric positive-definite (SPD) matrices.

Section 11 (Determinants I)

- What is a permutation and how to calculate its sign.
- What is a determinant, the Leibniz formula.
- Simplified formulas for 2×2 and 3×3 determinants.
- How to calculate determinants using Laplace (cofactor) expansion.
- How to calculate determinants using elementary row operations.
- How to calculate determinants of diagonal, block diagonal and triangular matrices.

Section 12 (Determinants II)

- About the determinants of elementary matrices and their inverses.
- That matrix A is invertible if and only if $\det(A) \neq 0$, and $\det(A^{-1}) = 1/\det(A)$.
- That for any two $n \times n$ matrices A and B it holds $\det(AB) = \det(A) \det(B)$.
- That $\det(A + B) \neq \det(A) + \det(B)$.
- That $\det(A) = \det(A^T)$ and what it means for column operations.
- How to use the so-called Cramer's rule to solve the linear system $Ax = b$.

- About the adjugate (adjunct) matrix and how to use it to calculate matrix inverse.
- About the relations between $\det(A)$, linear dependence of columns and rows in the matrix A , properties of the linear transformation $T(x) = Ax$, and the existence and uniqueness of solution to the linear system $Ax = b$.
- How to use determinants to calculate area and volume.

Section 13 (Linear Spaces)

- What is a linear space and how it differs from a standard set.
- How to check whether or not a set is a linear space.
- That linear spaces may contain other types of elements besides vectors.
- That a linear span always is a linear space.
- How to check if a set of vectors generates a given linear space.
- How to work with linear spaces of matrices.
- How to work with polynomial spaces.

Section 14 (Basis, Coordinates, Subspaces)

- What is a basis, and that a linear space can have many bases.
- How to find the coordinates of an element relative to a given basis.
- How to determine the dimension of a linear space.
- What is a subspace of a linear space, and how to recognize one.
- About the union and intersection of subspaces.
- How the change of basis in a linear space affects the coordinates of its elements.

Section 15 (Null Space, Column Space, Rank)

- What is the null space, how to check if a vector belongs to the null space.
- How to determine the dimension of the null space and find its basis.
- How is the null space related to the uniqueness of solution to the linear system $Ax=b$
- What is the column space, how to check if a vector belongs to the column space.
- How to determine the dimension of the column space and find its basis.
- How is the column space related to the existence of solution to the linear system $Ax = b$.
- How is the null space related to the kernel of the linear transformation $T(x) = Ax$.
- How is the column space related to the range of the linear transformation $T(x) = Ax$.
- What is the rank of a matrix, and its implications for the existence and uniqueness of solution to the linear system $Ax = b$.
- What does it mean for a matrix to have full rank or be rank-deficient.
- The Rank Theorem.

Section 16 (Eigenproblems I)

- What is an eigenproblem.

- About important applications of eigenproblems in various areas of science and engineering.
- How to verify if a given vector is an eigenvector, and if a given value is an eigenvalue.
- How to calculate eigenvalues and eigenvectors.
- About the characteristic polynomial and characteristic equation.
- How to handle matrices with repeated eigenvalues.
- About algebraic and geometric multiplicity of eigenvalues.
- How to determine algebraic and geometric multiplicity of eigenvalues.
- How to find a basis of an eigenspace.

Section 17 (Eigenproblems II)

- The geometrical meaning of eigenvalues and eigenvectors in R^2 and R^3 .
- That a matrix is singular if and only if it has a zero eigenvalue.
- That the null space of a matrix is the eigenspace corresponding to the zero eigenvalue.
- That eigenvectors corresponding to different eigenvalues are linearly independent.
- The Cayley–Hamilton theorem (CHT).
- How to use the CHT to efficiently calculate matrix inverse.
- How to use the CHT to efficiently calculate matrix powers and matrix exponential.
- About similar matrices, diagonalizable matrices, and eigenvector basis.
- How to use the eigenvector basis to efficiently calculate arbitrary matrix functions including matrix powers, the inverse matrix, matrix exponential, the square root of a matrix, etc.

Section 18 (Complex Linear Systems)

- About complex numbers and basic complex arithmetic.
- How to solve complex linear equations with real and complex coefficients.
- How to solve complex linear systems with real and complex matrices.
- How to perform complex elementary row operations.
- How to determine the rank of a complex matrix.
- How to check if a complex matrix is nonsingular.
- How to use the Cramer’s rule for complex matrices.
- How to check if complex vectors are linearly independent.
- How to invert complex matrices.
- How to transform complex linear systems into real ones.
- About the structure of complex roots of real-valued polynomials.
- About complex eigenvalues and eigenvectors of real matrices.
- How to diagonalize matrices with complex eigenvalues.

Section 19 (Normed Spaces, Inner Product Spaces)

- Important properties of the Euclidean norm and dot product in R^n .
- The Cauchy–Schwarz inequality and triangle inequality.

- How are norm and inner product defined and used in general linear spaces.
- About important norms and inner products in spaces of matrices and polynomials.
- How to calculate norms, distances and angles of matrices and polynomials.
- That every inner product induces a norm, but not every norm induces an inner product.
- About the parallelogram law.
- About the norm and inner product of complex vectors in C^n .
- About the norm and inner product in general complex linear spaces.

Section 20 (Orthogonality and Best Approximation)

- About orthogonal complements and orthogonal subspaces.
- How to find a basis in an orthogonal complement.
- That $\text{Row}(A)$ and $\text{Nul}(A)$ of a $m \times n$ matrix A are orthogonal complements in R^n .
- About orthogonal sets and orthogonal bases.
- About orthogonal decompositions and orthogonal projections on subspaces.
- That orthogonal projection operators are idempotent.
- How to calculate orthogonal decomposition with and without orthogonal basis.
- How to calculate the best approximation of an element in a subspace.
- How to calculate the distance of an element from a subspace.
- The Gram-Schmidt process.
- How to orthogonalize vectors, matrices, and polynomials.
- About the Fourier series expansion of periodic functions.
- How to obtain Legendre polynomials.
- How to use Legendre polynomials for best polynomial approximation of functions.

Section 21 (Spectral Theorem)

- About the basis and dimension of the complex vector space C^n .
- What are conjugate-transpose matrices and how to work with them.
- About Hermitian, orthogonal, and unitary matrices.
- That the eigenvalues of real symmetric and Hermitian matrices are real.
- That the eigenvectors of real symmetric and Hermitian matrices can be used to create an orthogonal basis in R^n and C^n .
- About orthogonal diagonalization of real symmetric matrices.
- About unitary diagonalization of Hermitian matrices.
- The Spectral Theorem for real symmetric and Hermitian matrices.
- About the outer product of vectors.
- How to perform spectral decomposition of real symmetric and Hermitian matrices.
- How eigenvalues are related to definiteness of real symmetric and Hermitian matrices.

- How to calculate eigenvalues of large matrices using Numpy.

Section 22 (QR Factorization and Least-Squares Problems)

- How to perform the QR factorization of rectangular and square matrices.
- What are Least-Squares problems and how to solve them.
- Properties of the normal equation, existence and uniqueness of solution.
- How to fit lines, polynomials, and bivariate polynomials to data.
- How to fit implicitly given curves to data.
- How to solve Least-Square problems via QR factorization.

Section 23 (Singular Value Decomposition)

- The basic idea of SVD.
- About the shared eigenvalues of the matrices $A^T A$ and AA^T .
- What are singular values and how to calculate them efficiently.
- That the number of nonzero singular values equals the rank of the matrix.
- How to calculate right and left singular vectors, and perform SVD.
- The difference between full and compact SVD.
- About the Moore-Penrose pseudoinverse and its basic properties.
- What types of problems can be solved using the M-P pseudoinverse.
- How to calculate the M-P pseudoinverse using SVD.
- How to calculate the SVD of complex matrices.
- How are singular values related to the Frobenius norm.
- What is the spectral (operator) norm of a matrix.
- How to calculate the error caused by truncating the SVD.
- How to use SVD for rank estimation in large data sets.
- How SVD is used in image processing.

Section 24 (Large Linear Systems)

- The condition number and how to calculate it using singular values.
- The condition number of real symmetric and Hermitian matrices.
- The role the condition number plays in the solution of linear systems.
- Mantissa, exponent, and the representation (round-off) error.
- Machine epsilon and finite computer arithmetic.
- Instability of the Gauss elimination.
- Memory issues related to storing large matrices.
- COO, CSR and CSC representation of sparse matrices.
- Sparsity-preserving and sparsity-breaking matrix operations.

Section 25 (Linear Algebra with Python)

- Import Numpy and Scipy.
- Define (large) vectors and matrices.
- Perform standard vector and matrix operations.
- Edit matrices and vectors using the Python for-loop.
- Extract row and column vectors from matrices.
- Use direct and iterative matrix solvers.
- Determine the rank of a matrix.

- Perform LU and Cholesky factorizations.
- Calculate determinants, eigenvalues and eigenvectors.
- Diagonalize matrices, calculate functions of matrices.
- Perform QR factorization and orthogonalize sets of vectors.
- Solve Least-Squares problems.
- Perform spectral decomposition and SVD.

Acknowledgment

The author thanks the University of Nevada, Reno administration for their supportive attitude towards the novel self-paced courses developed by NCLab.

References

- [1] ADA: Information and Technical Assistance on the Americans with Disabilities Act. URL https://www.ada.gov/2010ADASTandards_index.htm. Accessed September 20, 2020.
- [2] Chegg: Find Solutions For Your Homework. <http://chegg.com>. Accessed September 20, 2020.
- [3] G. Akçayır, M. Akçayır, *The Flipped Classroom: A Review of Its Advantages and Challenges*. Computers and Education **126** (2020), 334–345.
- [4] HTML5, URL <https://en.wikipedia.org/wiki/HTML5>. Accessed September 20, 2020.
- [5] Indeed.com: 7 Reasons Why it's So Hard to Get a Job After College, <https://www.indeed.com/career-advice/finding-a-job/why-is-it-so-hard-to-get-a-job-after-college>. Accessed September 20, 2020.
- [6] MacroTrends, Chegg Company Valuation <https://www.macrotrends.net/stocks/charts/CHGG/chegg/net-worth/>. Accessed September 20, 2020.
- [7] Mathjax: Beautiful and Accessible Math in All Web Browsers. URL <https://www.mathjax.org/>. Accessed September 20, 2020.
- [8] NCLab (Network Computing Laboratory), URL <http://nclab.com>. Accessed September 20, 2020.
- [9] NCLab's Self-Paced Linear Algebra Course URL <http://nclab.com/courses/linear-algebra>. Accessed September 20, 2020.
- [10] Python, URL <https://www.python.org/>. Accessed September 20, 2020.
- [11] A. Sams, J. Bergmann, *Flip Your Students' Learning*. Educational Leadership **70** (2013), 16–20.
- [12] P. Solin, *Using the NCLab's Matrix App to Obtain the Reduced Echelon Form of a 4×5 Matrix*. URL <https://youtu.be/2DjUTkqK5Fk>. Accessed September 20, 2020.
- [13] P. Solin, *Using the NCLab Matrix App*. URL <http://femhub.com/pavel/work/LinearAlgebraApp.pdf>. Accessed September 20, 2020.
- [14] TinyMCE: The World's Most Customizable and Flexible Rich Text Editor. URL <https://www.tiny.cloud/>. Accessed September 20, 2020.

Pavel Solin

Founder and Director, NCLab (<http://nclab.com>)

Professor, Department of Mathematics and Statistics, University of Nevada, Reno

e-mail: pavel@nclab.com, solin@unr.edu