Crossroads The ACM Magazine for Students

WINTER 2013 VOL.20 · NO.2

XRDS.ACM.ORG

XRDS

Getting Dressed in Tech

The Latest in Wearable Tech

Ori Inbar on Making Augmented Reality a Reality

How to Ace Google's Technical Interview

Association for Computing Machinery

tvx 2014

ACM INTERNATIONAL CONFERENCE ON INTERACTIVE EXPERIENCES FOR TELEVISION AND ONLINE VIDEO

25-27 JUNE, 2014
NEWCASTLE UPON TYNE UK

Paper Submissions by
3 February 2014

Workshop, Demo, WIP
DC, Grand Challenge
& Industrial
Submissions by
31 March 2014

Welcoming Submissions on

Content Production
Systems & Infrastructures
Devices & Interaction Techniques
Experience Design & Evaluation
Media Studies
Data Science & Recommendations
Business Models & Marketing
Innovative Concepts & Media Art

TVX2014.COM

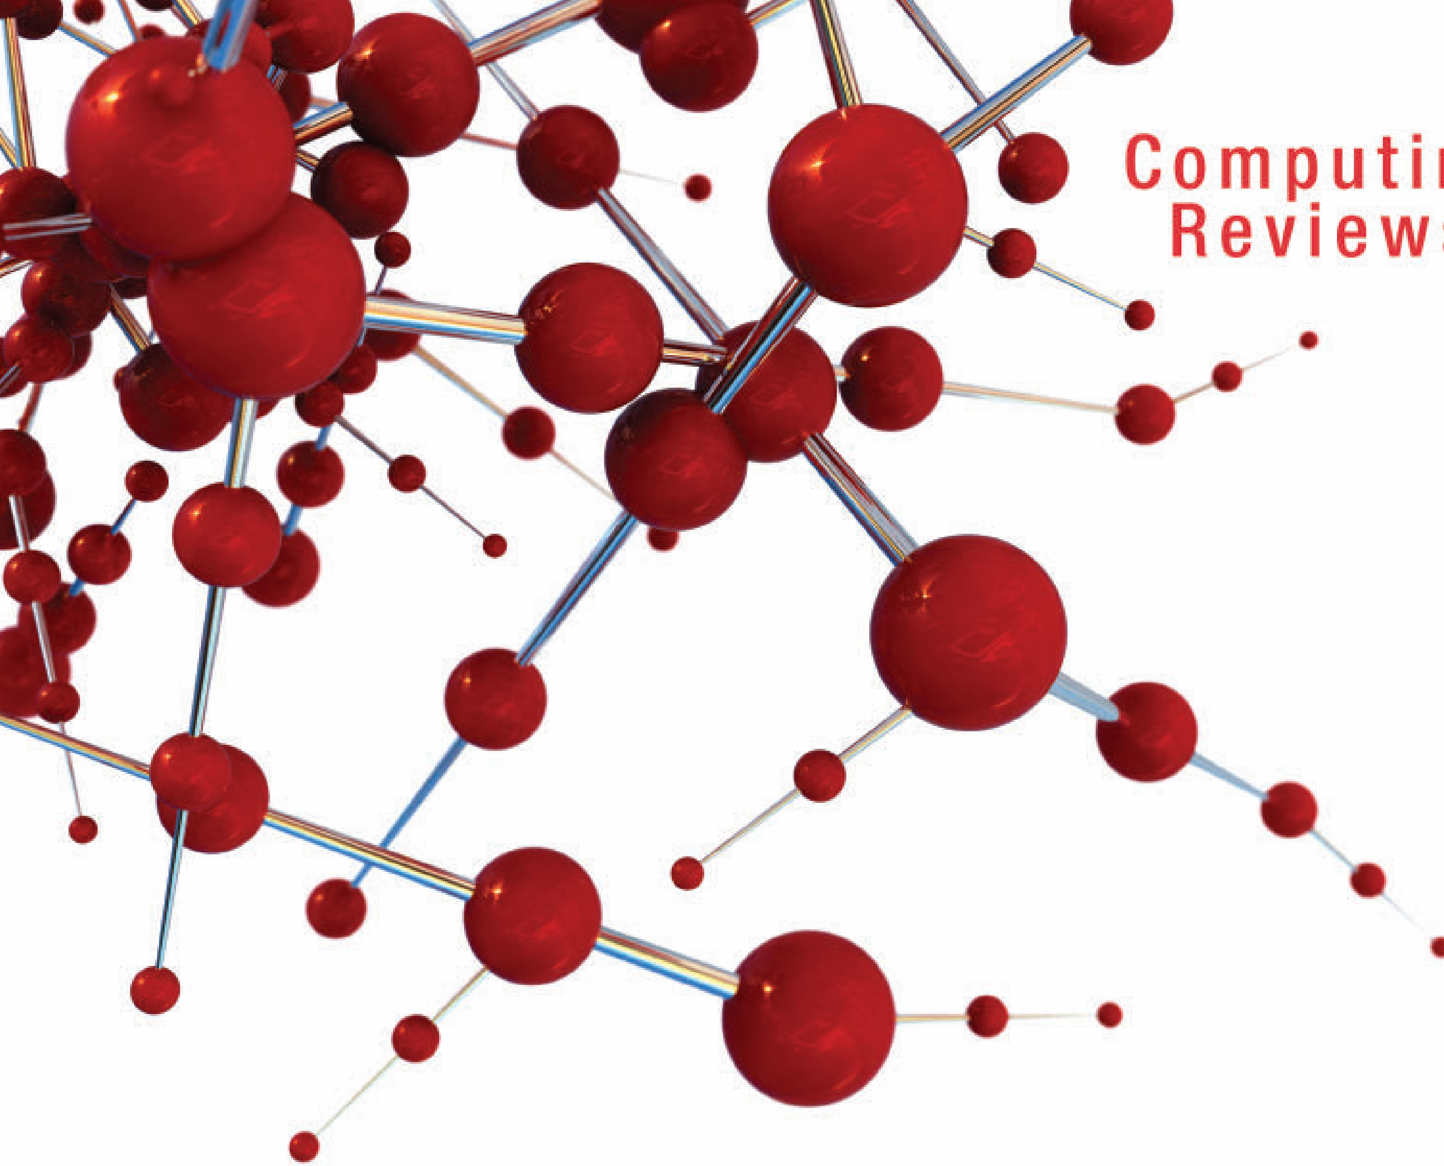

XRDS
Crossroads
The ACM Magazine for Students
WINTER 2013 VOL.20 · NO.2

Cover Illustration by Patrick George

Photograph by Ekkehard Streit

begin

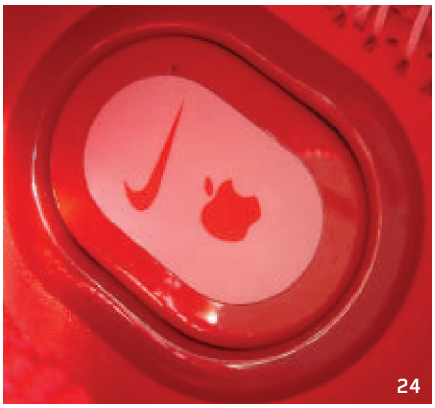

features

end

Photograph by Brandon Shigeta; Photograph by Volker Steeger

Association for Computing Machinery

XRDS

ACM
Association for Computing Machinery
2 Penn Plaza,
Suite 701
New York, NY
10121-0701 USA
+1 212-869-7440

CONTACT
General feedback:
xrds@acm.org

For submission guidelines, please see http://xrds.acm.org/authorguidelines.cfm

SUBSCRIBE
Subscriptions ($19 per year includes *XRDS* electronic subscription) are available by becoming an ACM Student Member www.acm.org/membership/student

Non-member subscriptions: $80 per year http://store.acm.org/acmstore

ACM Member Services
To renew your ACM membership or *XRDS* subscription, please send a letter with your name, address, member number and payment to:

ACM General Post Office
P.O. Box 30777
New York, NY
10087-0777 USA

Postal Information
XRDS (ISSN# 1528-4981) is published quarterly in spring, winter, summer and fall by Association for Computing Machinery, 2 Penn Plaza, Suite 701, New York, NY 10121. Application to mail at Periodical Postage rates is paid at New York, NY and additional mailing offices.

POSTMASTER: Send addresses change to: *XRDS: Crossroads,* Association for Computing Machinery, 2 Penn Plaza, Suite 701, New York, NY 10121.

Offering# XRDS0171
ISSN# 1528-4972 (print)
ISSN# 1528-4980 (electronic)

Forget About Blenders

We've been interviewing here at *XRDS*. As a student publication, our editors invariably graduate at some point, and along with their lab mates and thesis research they leave behind editing *XRDS*. Saying goodbye to a graduating editor is not easy; for starters, it always comes as a surprise: "What! X is already graduating?" (Of course, at this point X's graduation date has only been known for the past five years.) Then slight panic arises: "How are we ever going to find someone as good as X?"

There's also a personal attachment—you know the amount of thought and effort X put into making *XRDS* a great magazine. You owe her for that one time she came to your rescue right before the deadline. And there was the time the both of you shared a couple of beers and some good laughs during the last face-to-face meeting in New York (and maybe you even almost got arrested together…long story). In short, X has become an organic part of the team and you're really sorry to see her go. Then along comes new editor Y, and while he's not quite X, he brings fantastic new ideas and a fresh perspective. What seems like a mere couple of issues flash by and then, what? Y's graduation is next year already? How did we ever manage without him?

The interview process at *XRDS* is relatively straightforward, involving initial screening based on a candidate's CV and online material (personal webpage, etc.), an interview to check whether there's a good mutual match, and then a written edit test and evaluation. We try to recruit students from diverse backgrounds, and it's very important to us to get the team right—as anyone who's worked as part of a team knows, a single unproductive member can result in everybody dragging their feet. Luckily, the other side happens too: One enthusiastic and energetic member can drive the rest forward. To a large extent, it's all about motivation.

At smaller companies, given that you have the right skills and background, passing an interview can sometimes be a matter of chemistry with whoever is interviewing you (we know of at least one case where an interviewee was accepted after bonding with his interviewer over a shared appreciation for a certain book, after spending most of the interview talking about it). This is not the case when interviewing with any of the Silicon Valley giants. At large companies like Google, Apple, or Facebook, the recruiting process is much more complex.

In this issue we have a fascinating insider's look at what a Google interview is all about. This is a process you cannot charm your way through or rely

UPCOMING ISSUES

Spring 2014
[March issue]
Air, Water, Land, and Technology-Ubiquitous Earth

Summer 2014
[June issue]
Diversity in Computer Science
Article deadline: March 14, 2014

on your quick wit or luck—in fact it's designed to avoid precisely that. While you definitely will need all of the above to some extent, at least at Google and other large companies it's more like taking the SAT. Forget about blenders, you need to study, prepare, and relentlessly practice.[1]

Tech startups, however tend to take a different approach to the interview process. Referrals from current employees are vitally important. Even for the best candidates, it can be tough to get an interview with an exciting startup without knowing someone who's working with the existing team. One of the very best ways to set yourself apart from the rest is to build something relevant to the company before you submit your application, and point it out as quickly as you can in that first email or submission. Having code on github helps, even if it's not widely used. It gives the engineers an opportunity to evaluate your skill set without having you jump through hoops. The best startups will also be on the look out for a strong fit with their culture. Being able to get along well with your coworkers is a lot more important at a company of 10 or 15 than it is at a place with 3,000 where you could conceivably just transfer to another team.

Although the Web is full of very good advice on how to prepare for interviews, here's one tip that's maybe less common: Try taking the other side of the table. When faced with the need to decide among a few candidates in a limited amount of time, you realize the crucial things to get right during an interview. If you don't have an opportunity to take part in conducting a real interview, do a mock one with friends, taking turns as interviewer/interviewee. You can also do this during an actual interview, though it requires some delicacy and care. Ensuring the position and the company are a strong fit for you is just as important as the other side of the equation.

One enthusiastic and energetic member can drive the rest forward. To a large extent, it's all about motivation.

If the thought of all that work preparing for industry interviews (whether for an internship or full-time position) has not been stressful enough already, Chand John's recent blog post for the *Chronicle of Higher Education* is a must-read for graduate students. His discussion on the Ph.D. industry gap is a good reminder to get industry experience early on—if possible—for those considering moving on to industry after graduation.

We'd like to take this opportunity to send a warm farewell and good luck wishes to graduating XRDianS: Hannah (who masterfully led the issue you're holding), Debarka, John, and Luigi. Please accept our heartfelt thanks for your invaluable contributions; and a warm welcome to new members of the team—Hanieh, Virginie, Apoorvaa, and Bryan.

—Inbal Talgam-Cohen and Peter Kinnaird

P.S. This will be my last issue as co-EIC of XRDS. *After mulling things over for at least a year, I decided to leave academia to join a startup called Crowdtilt after five and a half years of post-graduate training. Working with Inbal and the rest of the* XRDS *team has been an incredible opportunity. Thanks to everyone who has contributed to* XRDS, *authors and editorial staff alike. We've made a great magazine together!*

—Peter

P.P.S. I'd like to add special thanks to Peter, who has led and shaped this magazine as co-editor-in-chief for the past two years. Your passion, creativity and leadership will be greatly missed. Good luck!

—Inbal

1 Google famously posed the following question to interviewees: "You are shrunk to the height of a nickel and thrown into a blender. Your mass is reduced so that your density is the same as usual. The blades start moving in 60 seconds. What do you do?" You can read the entire aritcle, "How to Ace a Google Interview" by William Poundstone, which ran in *The Wall Street Journal*. .

begin

INBOX

READ ALL ABOUT IT!

What is public and private anyway? See my essay in the #ACM student magazine #XRDS for a pragmatist answer: http://goo.gl/NZPSqA
—**Andreas Birkbak**, *Ph.D. Fellow, Aalborg University Copenhagen, Twitter (@communaut)*

XRDS special issue on #privacy and #anonymisation, lots of relevant articles http://tinyurl.com/p4pzlcf
—**UK Anonymisation**, *Twitter (@ukan_net)*

XRDS: Crossroads, The ACM Magazine for Students - The Complexities of Privacy and Anonymity http://bit.ly/1c1zTAv The Complexities of...
—**Little ODI Robot**, *Twitter (@LittleODIRobot)*

A MUSICAL NOTE

Thanks to @XRDS_ACM A fascinating article on "The well-programmed clavier: style in computer music composition" http://xrds.acm.org/article.cfm?aid=2460444
—**Zen Loves Sarmistha**, *Research Scholar/ Philosopher/Computer Scientist, Twitter (@ SarmisthaIsLove)*

@BitterRancor glad you enjoyed it Arun! I edited that one.
—**Ryan Kelly**, *Ph.D. student, University of Bath, ACM XRDS feature editor, Twitter (@rhyan2438529)*

@rhyan24385 Oh, very good, Ryan. Yes, I enjoyed it. The articles has n important discovery about clavier music composition.
—*@SarmisthaIsLove*

@rhyan24385 I shall pass on this article to my philosopher friend, who is a musician esp. a performer
—*@SarmisthaIsLove*

@BitterRancor great!
—*@rhyan2438529*

How to contact *XRDS*: Send a letter to the editors or other feedback by email (xrds@acm.org), on Facebook by posting on our group page (http://tinyurl.com/XRDS-Facebook), via Twitter by using #xrds in any message, or by post to ACM Attn: XRDS, 2 Penn Plaza, Suite 701, New York, New York 10121, U.S.

Photograph by Ekkehard Streit

INIT

Getting Dressed in Tech

"Wearable Computing" refers to embedding sensors and computation devices on the body in a seamless, unobtrusive, and invisible way.

But the trend toward wearable computing is not as new as it seems. Taking something useful and making it more convenient is as old as time. This can be seen with the historical shift from pocket watches to wristwatches. We're now seeing new tools and technologies being taken out of our pockets and put into seemingly more convenient places like our wrists, feet, and faces—think Google Glass.

These trends might revolutionize the way we live, behave, and interact. However, we know convenience and wow factor won't be the only driver in the explosion of wearable devices. Beyond the current commercial applications, extensive research is being done to push the limits of wearable computing. Shoes, watches, gloves, hats, and other common body-worn items could become a dynamic sensor network, providing vast quantities of information about the wearer and their surroundings.

Providing a comprehensive picture of recent trends in wearable computing seems infeasible. Thus, in this issue we highlight applications in personal behavior monitoring, health care, and human-computer interaction.

PERSONAL BEHAVIOR MONITORING

Sports are an everyday activity for millions of people. But not everybody has access to a trainer. Prompting Christina Strohrmann and Gerhard Tröster to introduce a wearable sensor system to monitor a runner's performance in the field. They suggest methods that assess skill level, technique, and fatigue from motion sensor data. For those that are less active, fitness trackers can be a motivator to get moving. Andrew Miller shares his experience using Fitbits and social media concepts with teens to help increase activity and stave off obesity.

An active body is ideal, but we mustn't forget about mental agility. Reading is a highly frequent activity that trains our mental fitness. Kai Kunze proposes a novel method to recognize document types based on gaze tracking. Integrating such a system into glasses could improve awareness about our personal knowledge acquisition in the future.

HEALTHCARE APPLICATIONS

Wearable computing opens up a number of possibilities in healthcare. Current research experiences shift from acute care toward long-term monitoring with a focus on outpatient prevention and rehabilitation. Gabriele Spina and Oliver Amft suggests a smartphone-based training system that successfully assesses the performance and execution quality of exercises in rehabilitation patients. The model offers personalized training and provides feedback to patients at home.

As part of patients' recovery treatment, Rolf Adelsberger developed a wearable sensor system combining a pressure sensor sole and motion sensors to automatically assess posture stability. Mladen Milosevic, Aleksandar Milenkovic, and Emil Jovanov cover concepts on physiological monitoring systems and monitoring applications using standard smartphone sensors as part of an mHealth project.

HUMAN-COMPUTER INTERACTION

In the age of smartphones, tablets, and public screens interaction with computers is a common occurrence. Almost everyone has experienced situations where the operation of tiny screens and keys can be challenging. Christina Amma and Tanja Schultz introduce a glove that could replace the key operation with air-writing. Rounding out the isuse, is Viswam Nathan's discussion on the challenges of creating a wearable EEG system, which can be used for a better brain-computer interface. This work could lead to touch-less input for many devices.

The future of wearable technology is much more than entertainment, convenience, and apps. The concepts that researchers are working on will definitely create a world of new opportunities. We're hopeful that this issue gives you a good sense about what is currently happening with wearable devices research, what the future could be, and spark some ideas on how you can add to it.

—Terrell R. Bennett and Julia Seiter, Issue Editors

The Nest thermostat, which went on sale in 2011, can be controlled over the Internet.

Fitbit, a Bluetooth pedometer and sleep sensor, raised a $12.5 million funding round in 2012.

BENEFIT

Student Chapters in Europe

ACM has 33 European student chapters, which until recently, have felt relatively isolated. This is why, in 2012, the ACM Europe Council decided to create a Council of European Chapter Leaders (CECL), which aims to facilitate chapter operation, increase chapter effectiveness, and serve as a link between chapters and between ACM and chapters.

A major aim of CECL is to create a community of European chapters. The CECL Facebook group is set up for chapter officers to engage in discussion, submit ideas, and share experiences. Chapters can create collaborative relationships and increase visibility, with CECL providing an easy and fast channel to ACM.

Another aim is to increase European participation in the Distinguished Speakers Program (DSP)and the ACM International Collegiate Programming Contest; the oldest, largest, and most prestigious programming contest in the world. Students can nominate more Europeans to the list of supported speakers, here http://dsp.acm.org/nominate_form.cfm.

To find out more about CECL, visit http://europe.acm.org/chapters.html.

—*Virginia Grande*

ADVICE

How You Can Change the World

Students are busy. Yes, there's no doubt about it. Between coursework, research, group meetings, conferences, there's very little time left in the day. And in this controlled chaos, we sometimes lose sight of what we are a part of.

But take a moment to consider what those outside our field see: A field filled with nuances and complications; a field of dizzying proportions that is slowly being applied to every aspect of their lives. When you stop and think about the rapid-pace of computing, it's easy to see why many think computer science is not unlike the florescent green falling digits of "The Matrix."

Is this how we want computing to be viewed by the public, as inaccessible and labyrinthine? Absolutely not. We should make research and education equal parts in all of our work.

It is our responsibility to change the way the public sees computing. By showing our communities exactly why we choose to study computing, we can grow along with them. This kind of outreach has finally gone mainstream thanks to famous physical scientists like Brian Greene, Neil DeGrasse Tyson, and the late Carl Sagan. But other than projects like IBM's Deep Blue and Watson, computer science has fallen behind in the realm of scientific outreach. We, as students, can change that. We can build this outreach into our education, to not only teach our communities but also to enhance our own educations.

Outreach provides a venue unlike any professional conference. You can no longer use the jargon, shortcuts, and acronyms with which you've become so familiar. You can no longer assume your audience has an all-encompassing knowledge of your field. As you begin to adapt to this vastly different type of audience, you begin to form a new understanding of the field you thought you knew.

Think of community outreach as job training. You will get to practice how to lead a class, how to answer questions quickly and concisely, and, most importantly, start to become more comfortable in front of a large audiences. Through community outreach not only do you have the opportunity to inspire the next generation of students, but you also have an invaluable opportunity to practice the skills you will spend the entirety of your life using.

So why not get involved? See if your university has a science outreach program that you can join. Find out if your city is hosting a day of civic hacking or a hackathon. Help your department host a day for community members to come and have their computers diagnosed, repaired, and even properly recycled. Offer to teach a programming or Web design class at your local library. Volunteer for a local FIRST Robotics or Lego Mindstorms Team. The possibilities are endless. Get out there and share what you're passionate about with your community.

We have both the opportunity and the responsibility to show the world the true beauty and intricacy behind computing. We can be the difference.

Biography

Connor Bain is a junior at the University of South Carolina Honors College studying computer science, mathematics, and music. He also serves as the Director of Carolina Science Outreach (www.csousc.org). Upon graduation, he plans to attend graduate school in computer science and eventually enter academia.

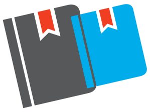

RFID tags in library books and library cards allow libraries to detect which books are stolen and who is stealing them.

UPDATES

Maintaining ACM Traditions

Professional Development Done Right

As ACM student chapters grow and reach out to new members, chapter leaders are faced with the challenge of trying to host activities that spark varied interests. Typically our issues feature student chapters that were successful in pushing boundaries and offering unique outlets for student passion and creativity. Chapters around the world are accomplishing incredible feats, which we choose not to overlook. However, in this issue of *XRDS* we are highlighting the tradition behind ACM student chapters—advancing computing as a science and a profession while enabling professional development.

When it comes to keeping up traditions, look no further than the ACM student chapter at Florida State University (FSU). Formed in 1990, this student chapter has seen its fair share of involved undergraduate and graduate members transform into active faculty members who seek to advise future generations of ACM members at FSU. As with most ACM student chapters, comes the traditional programming competition, which has become a flagship event for the ACM student chapter at FSU. This event alone garners roughly 30 to 40 teams of one to three students per team. Following the success of this local competition, students are then coached and practice with professors and graduate students for participation in regional programming contests. These are rewarding events for all students, regardless of skillset or age.

The FSU chapter holds its technical workshops and professional development in high regard. Student leaders work hard to organize various workshops on new technologies, honing software development skills and preparation for industry employment. The first series this fall at FSU was on Git/Mercurial and best practices for software versioning. Frank Valcarcel, president of the ACM student chapter at FSU said the following about the series: "This is a skill we feel is very important for today's job candidates to have and we don't touch on it much in our curriculum. We want to expose our members to it as best we can. We do this by applying the principles to what the workflow for an average student would be. Things like branching, merging, and tagging would be useful when a student is working on an assignment in a programming class, and we feel that by tailoring the workshop to that will help convey the power and versatility and encourage our members to continue on with it after the workshop."

ACM Members at a recent "Hacking The Interview" workshop, earlier this semester.

Using the momentum behind their technical workshops, the chapter leaders also invite local and regional companies to hold talks with its members. Valcarcel elaborated on these efforts: "These talks help prepare them mentally for the tough application and interview process associated with jobs in industry. This is something we are improving upon this year. Our chapter has begun planning a series of mock interviews and interview workshops to take our next round of grads to the top. We have asked representatives and colleagues to provide sample questions and answers to help make the experience as authentic as possible. We will begin each of these events with a game of computer science trivia. Where students will team up and then answer questions in Jeopardy-like fashion. We do this to help the students familiarize themselves with industry terminology and become more confident with their answers."

It is clear the members of the ACM chapter at FSU are a passionate group seeking to advance computing and the knowledge of their peers. Activities within this chapter are not only thoughtful and purposeful, but also aligned to best meet students where they stand academically, as well as where they will be moving professionally. If you would like to learn more about the ACM activities at FSU you can visit their website at: http://fsu.acm.org.

—Michael Zuba

$10 M
The Pebble smartwatch raised more than 100 times its funding goal on Kickstarter.

CAREERS

The Google Technical Interview
How to Get Your Dream Job

Most successful students would not consider taking a final exam without preparation. For students—undergraduate and graduate—about to head into industry, job interviews benefit from the same level of consideration. There are many books and articles on interview skills, and most academic career centers offer additional training. What this article covers is more specific: A highlevel view of Google's engineering interviews.

The interview process at Google has been designed (and redesigned!) from the ground up to avoid false positives. We want to avoid making offers to candidates who would not be successful at Google. (The cost of this unfortunately includes more false negatives, which are times when we turn down somebody who would have done well.) The recruiters and engineers you will speak with want to see where you shine, whether you can do the job, and make sure you're someone they want to work with. This article is designed to help both you and Google achieve those goals—and help the interview be an interesting, even pleasant, experience, too.

You will meet at least two types of Googlers (Google employees) in the interview process. The first are our recruiters. Recruiters are nontechnical employees who are experts at both finding candidates and helping them through the interview process. The second are our technical interviewers; they are fulltime engineers who volunteer to help with the hiring process by interviewing candidates like you. All of our Googlers come from academic backgrounds and from industry, and can answer most (if not all) of the questions you might have, including whether Google is likely to be a good fit for you. So please ask away!

There is a standard format for most technical interviews. (Ph.D. students and more experienced candidates may be given one additional interview with a slightly different format, but similar advice applies.) For about 45 minutes you meet with a single technical interviewer, who will present a programming problem and ask you to work out one or more solutions to it. In some interviews, you will be asked to code up one of your solutions on a whiteboard. All of our questions have multiple solutions, and some of our questions do not have a single best answer, so if you have more than one solution, explain the tradeoffs or the benefits of your preferred solution.

Each interview day will have up to five of these 45-minute interviews, depending on your schedule, proximity to the nearest office, and interviewer availability. Candidates living farther away may start with phone interviews before proceeding to an onsite interview. The types of interviewers and which questions they ask are the same in both cases.

Overall, the interviewers are simply trying to decide one thing: Would you be a good fit for Google?

Let's break down a typical interview, piece by piece.

The programming problems are not "trick" questions, but they will always have aspects that require care and attention. First, make sure you understand the problem properly. It helps to clarify assumptions before diving in too deeply, and if you are confused about the question, do ask for examples, or for the question to be reworded. The interviewer will often not offer information until you ask for it. "How big could the input be?" "What happens with bad inputs?" and "How often will we run this?" are three common clarifications. If you're stuck, one thing to do is recheck your assumptions.

After you have the clarifications that you need, dive into solving the problem. There are usually several paths to a good solution. Much like math homework, it is essential to show your work; the way to do this is to talk to the interviewer, and explain what you are thinking. This is easy for some of us, but really hard for others, so it is important to practice this skill. If narrating your entire thought process will significantly reduce your ability to think on your feet, it is OK to be quiet for a minute or two—but then tell the interviewer what you considered, and why you chose what you chose. The more you can communicate your thought process, the better. If this might be hard for you to do, it is definitely worth practicing with a friend.

Those in the Quantified Self movement collect data about themselves with sensors,with the goal of making better life decisions.

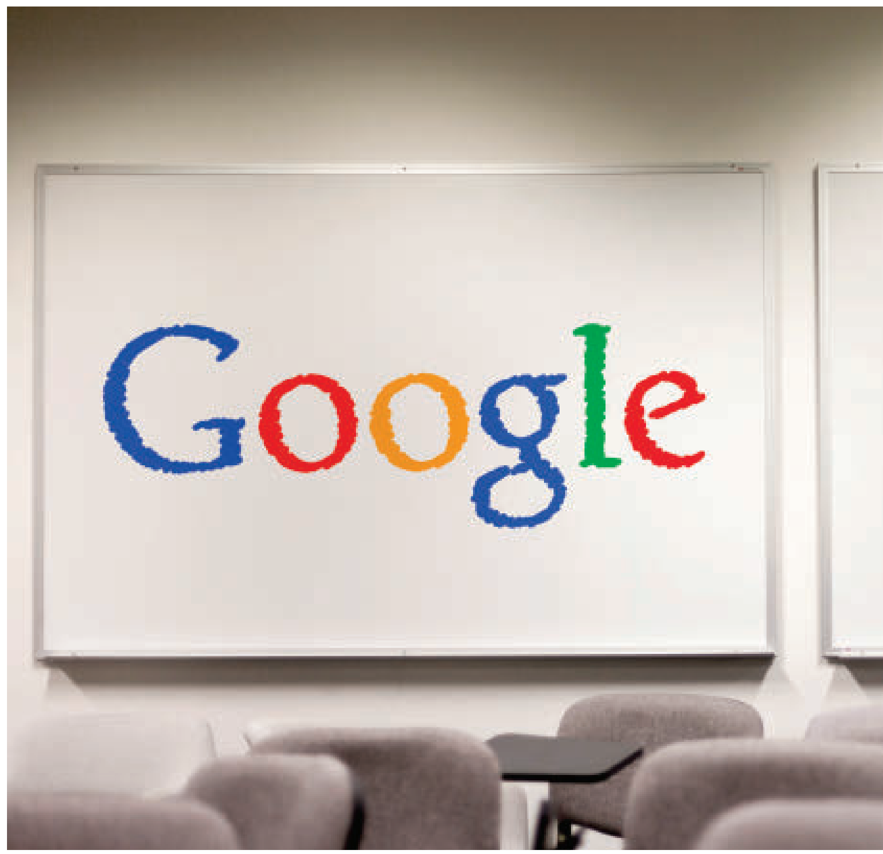

Feel free to give answers you know are imperfect; explain briefly why they are not the best answer, and keep going. The first solution is rarely the best, and a sequence of answers very clearly shows your thought process toward solving the problem. It is also fine to start with a brute-force approach, as it gives an initial benchmark for the answers that follow. One trap to avoid is getting stuck thinking about incremental improvements to the worst algorithm; sometimes you will need to leap to another approach.

One of the most important things you should know is Big O notation and the analysis it represents. It is the common language for discussing algorithmic performance. While Big O is normally used to discuss how well an algorithm will scale, it is also good to consider disk, memory, network, and other needs for each solution.

At any point in an interview it is fine to ask if you can move to the whiteboard to take notes, draw diagrams, or explain what you are thinking. You are also welcome to use pen and paper if it will better help you keep your thoughts organized. In many interviews, the interviewer will ask you to move to the whiteboard, and then it is usually time to write some code.

For our engineering positions, you will need to know a language like Java, C++, or Python. Knowing more than one is a nice touch, but you must know one well. For each interview with a coding component, you will usually write between 10 and 50 lines of code.

If you are rusty at coding, you should practice. If you code every day, you should still prepare by solving a few interview-like questions. Candidates who prepare do better. You should also practice writing code on a whiteboard, and explaining what you are doing while you do it. A whiteboard is significantly different than using your favorite editor or IDE. It also helps to have a friend or mentor review your practice code. As you know, code that works perfectly but is very hard for others to read is usually not a good idea in team-based software development. We are not worried about handwriting, but we do actively look for clean, maintainable code.

Testing code is important when working as an engineer, and it helps in interviews, too. After writing code, you should test it. Run a normal input, try the edge cases, and see if what you wrote behaves properly. This helps many candidates bump a mediocre answer into a significantly above-the-bar performance. When I interviewed for my position at Google, this was the difference between code that worked and a likely failure

If you are unsure about something, you should be open about your uncertainty. That will help you, not hurt you. We value honest communication. Likewise, if you are stuck, it is OK to ask for help if you need it—but fishing for answers is a bad tactic. If the interviewer gives you a hint, it's always a good idea to listen to it, and consider what they said: Their comments aren't random, and they are trying to help you. If you are working on an idea, and the interviewer is silent (or just taking notes), do not worry that something is going wrong. They are waiting for you, and

A smart electricity grid can adapt to prevent small failures from producing large blackouts.

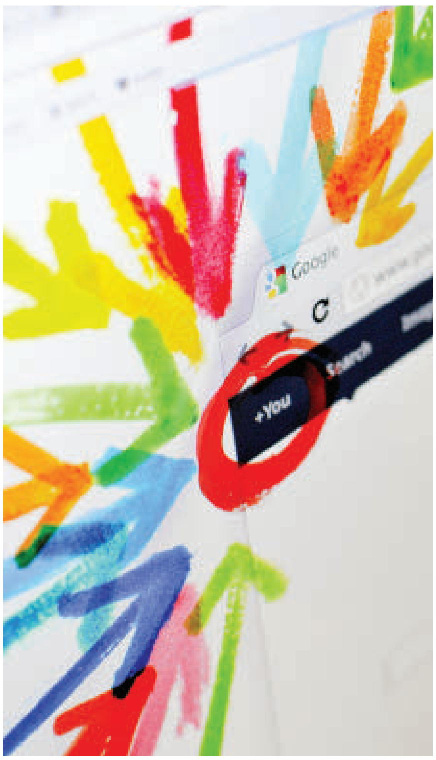

most importantly, they are finding it worth the wait.

Data structures and algorithms are the most important classes in the undergraduate curriculum. (Really!) If your school offers higher-level algorithms classes, take those as well. The majority of successful applicants took these courses, enjoyed them, and studied the material recently. Steve Yegge wrote about this a few years back in his blog post, "Get That Job At Google." I owe my job to reading his post; you should read it too.

Overall, the interviewers are simply trying to decide one thing: Would you be a good fit for Google? Determining that involves answering several other questions. Are you someone they want to work with? Are you someone who would make their team better? Are you someone they want writing code they will use and depend on? Can you think on your feet? Can you explain your ideas to coworkers? Can you write and test code? And are you friendly enough to chat with every day?

At the end of each interview, there will usually be a few minutes for you to ask questions. Have a few ready! The interviewers can tell you what it is like to work here, what we love, what we hate, how often we travel (or don't!), and are open to answering pretty much anything you remember to ask. (With one exception: We can't tell you how you did.) This is a great time to make sure that Google is the right place for you.

A typical 45-minute interview will consist of 35 minutes of programming problems, five minutes of questions, which leaves only five minutes for everything else, including an introduction, discussing your resume, and asking questions about your prior work experience. That said, your resume and work experience do matter; that and any references are what get you into an interview, so please make sure to polish your resume. Every interviewer you will meet will have been given your resume several days in advance, which is part of what helps them choose their questions.

Each interviewer has a limited amount of time to convince themselves that you will be a great hire, and they want to spend that time in the most efficient way. Therefore once you are in a technical interview, our interviewers will mostly focus on programming problems, not the resume, which we find to be the best use of your time.

Although it's unlikely to be the focus of an interview slot, be prepared to discuss what is on your resume. You should be able to talk about your experiences (especially the technical bits), explain your areas of focus and why they interest you, and be able to describe your contributions to the projects you list. The typical interview questions also apply: Why do you want to work for Google, and which types of projects are the most interesting to you?

A question that comes up time and again is: What should I wear? For Google, the advice I hear repeatedly, and seems to hold true, is: "Wear something that makes you feel comfortable." The specifics are up to you. That said, I feel it is still worth spending a bit of time on grooming. This will be your first time meeting meeting people you may work with for years, and making a decent impression helps convince the interviewer that they want to work with you.

In summary: Refresh your knowledge of data structures, algorithms, and writing clean code on a whiteboard. Come to the interview well rested, and feel free to ask the recruiter questions ahead of time. Be able to talk about your experience, and be ready to spend most of your time on the programming problems. Once in the interview, feel free to ask questions about the problem you are working on. During the interview, make sure to "talk out loud" enough. When you make decisions on how to solve something, make sure the interviewer knows about it. And be sure to ask questions that will help you find out if Google is a good fit for you.

Finally, be who you are, and be the best version of yourself. Our recruiters liked you, and the odds are our engineers will like you, too. Good luck!

(Parts of this have appeared in "Anatomy of the Google Interview," a talk given at Google by Carl Evankovich. Without his work there, this article wouldn't have happened; thank you!)

Biography

Dean Jackson's a member of the ACM and an engineer working at Google Pittsburgh, focused on Google Ads, and a frequent contributor to Google's recruiting programs.

Photograph by Ahmad Faizal Yahya

The *XRDS* blog highlights a range of topics from security and privacy to neuroscience. Selected blog posts, edited for print, will be featured in every issue. Please visit xrds.acm.org/blog to read each post in its entirety. Keeping with our theme of professional development, included is a guest post on how to craft a publishable research paper.

BLOGS

Photograph by Mike Vadala

Security Bugs in Large Software Ecosystems

By Dimitris Mitropoulos

In a previous blog post, I discussed the occurrence of security bugs through software evolution. In this post we will examine their existence in a large software ecosystem. To achieve this—together with four other colleagues (Vasilios Karakoidas, Georgios Gousios, Panos Louridas and Diomidis Spinellis)—we used the FindBugs static analysis tool, to analyze all the projects that exist in the Maven central repository (approximately 260GB of interdependent project versions).

Let's address the more straightforward question first. What is the best way to present an algorithm? How descriptive and specific should it be? Should it be entirely self-contained or, for instance, could we have a pointer to a "... subroutine of choice?" Is implementability more important than readability?

Maven is a build automation tool used primarily for Java projects and is hosted by the Apache Software Foundation. It uses XML to describe the software project being built, its dependencies on other external modules, the build order, and required plug-ins. First, we scanned the Maven repository for appropriate JARs and created a list. After some project filtering, we narrowed down our data set to 17,505 projects with 115,214 versions. With the JAR list at hand, we created a series of processing tasks and added them to a task queue. Then we executed 25 (Unix-based) workers written in Python that checked out tasks from the queue, processed the data after invoking FindBugs, and stored the results to a data repository.

FindBugs separates software bugs into nine categories. Two of them involve security issues: security and malicious code. From the total number of releases, 4,353 of them contained at least one bug coming from the first category and 45,559 coming from the second.

Together with the bad practice bugs and the style bugs, security bugs (the sum of the security and malicious code categories) are the most popular in the repository (≥ 21 percent each). This could be a strong indication that programmers write code that implements the required functionality without considering its many security aspects, an issue that has already been reported in literature.

Another observation involves bugs that we could call "severe" and they are a subset of the security category. Such bugs are related to vulnerabilities that appear due to the lack of user-input validation and can lead to damaging attacks like SQL injection and Cross-Site Scripting. To exploit such vulnerabilities, a malicious user does not have to know anything about the application internals. For all the other bugs, another program should be written to incorporate references to mutable objects, access non-final fields, etc. Also, as bug descriptions indicate, if an application has such bugs, it might have more vulnerabilities than FindBugs reports. In essence, 5,501 releases (≈ 4.77 percent) contained at least one severe security bug. Given the fact that other projects include these versions as their dependencies, they are automatically rendered vulnerable if they use the code fragments that include the defects.

Linus's Law states, "given enough eyeballs, all bugs are shallow." In a context like this, we expect the project versions that are dependencies to many other projects would have a small number of security bugs. To examine this variation of Linus's Law and highlight the domino effect we did the following: During the experiment we retrieved the dependencies of every version. Based on this information we created a graph that represented the snapshot of the Maven repository. The nodes of the graph represented the versions and the vertices their dependencies. The graph contained 80,354 nodes. Obviously, the number does not correspond to the number of the total versions. This is because some versions did not contain any information about their dependencies so they are not represented in the graph. After creating the graph, we ran the PageRank algorithm on it and retrieved all PageRanks for each node. Then we examined the security bugs of the 50 most popular nodes based on their PageRank. Contrary to Linus's Law, 33 of them contained security bugs, while two of them contained severe bugs. Twenty-five of them were latest versions at the time. This also highlights the domino effect.

Future work could also involve the observation of other ecosystems, that serve different languages than Java, in the

99

The percentage of Americans who live in areas with cell phone coverage.

---

same manner such as, Python's PyPy (Python Package Index) and Perl's CPAN (Comprehensive Perl Archive Network).

Dimitris Mitropoulos is a Ph.D. candidate at the Athens University of Economics and Business. His research interests include information security and software engineering. During his studies he has worked on several research projects and is the author of a number of open-source software libraries.

The Scary Reality of Identity Theft

By Wolfgang Richter

One of the most basic philosophical questions stems from attempting to identify oneself, with the first step of proving you actually exist. René Descartes provides a proof with *Cogito ergo sum*, meaning "I think, therefore I am." The intuition is that the mere fact of thinking forms a proof that you exist. But who or what are you exactly? What identifies you? How can we definitively prove you are what you claim to be? Who you claim to be? The problem of identity is an incredibly hard one—how do you know a letter in the mail is from the person that signed it? How do you know a text was written by the owner of a certain phone? How do you know an email comes from the person that owns an email address? This is a fundamental problem that faces the fields of computer science and cryptography, and it is incredibly hard to solve.

We've all encountered spam emails or emails laden with viruses from friends whose accounts or computers became compromised. People have broken into the accounts of well-known celebrities online and sent messages in their name. News items featuring databases of personal information from compromised firms as large as Sony and the PlayStation Network occur with striking regularity. Protecting your identity is a losing battle, and in your lifetime it is almost guaranteed that your identity will be stolen at least once, if not multiple times.

What does identity theft look like? I don't want to directly post personal information directly, but if you click through the link I referenced on the blog you will see real identity theft from a sample left by criminals for people who buy credit cards. Yes, in the seedy underworld of the Internet, credit cards are a form of currency and they are traded in quantities of a thousand to millions at a time.

Faced with the certainty that your identity will be stolen, what can you do? Banks and credit card companies use statistics to identify questionable transactions and automatically freeze your debit or credit cards if they detect something odd. Clearly, there is an incentive to stop identity theft as early as possible. But that's what the big corporations do—they build computer systems to monitor transactions of all their customers. What can you do? The only thing you can do is to minimize the damage when an identity theft occurs. There are multiple methods of doing this, but here are three quick tips I follow:

1. **CCD.** Use cash over credit, and credit over debit
2. **Credit Reports.** Check your credit reports regularly throughout the year, it's free.
3. **Statements.** Check your credit and debit statements as often as possible

The first tip, CCD, is a proactive tip and the other two are reactive tips, but all of them should be used together to help minimize damage when identity theft occurs.

Let's look at CCD. The idea here is to limit exposure to your primary asset: your money. When someone steals your information, make sure they can't steal your money. If you only dealt in cash, you would never expose your identity. But that's infeasible in today's world of online transactions, which are predominantly card only. So, my recommendation is to use a credit card whenever you can't use cash. That way, when your information is stolen, a financial institution's money will be stolen—not your money. Credit cards create short-term loans between you and a financial institute. At least initially, your money is never involved. As a last resort, use a debit card when you have to. But, understand that when your information is stolen your bank account balance will drop to zero or go negative. You will no longer be able to pay bills, checks will bounce, and your life will be hell until you *hopefully* get the money back.

Wolfgang Richter is a fifth year Ph.D. student in Carnegie Mellon University's Computer Science Department. His research focus is in distributed systems and he works under Mahadev Satyanarayanan. His current research thread is in developing technologies leading to introspecting clouds.

Photograph by Stephen VanHorn

5.3

The average number of WiFi-connected devices each student has in their dorm room at one U.S. college.

Mobile processors are approaching the compute capability of last-generation gaming consoles.

The Many Stages of Writing a Paper, and How to Close the Deal

Originally posted on The Geomblog

By Suresh Venkatasubramanian

Producing a piece of research for publication has many stages, and each stage has different needs, requiring different ways of operating. Learning these stages is a key developmental step for a graduate student.

From my conversations with students (mine and others), I think this is how students think a paper gets written:

1. Advisor produces problem miraculously from thin air.
2. Come up with solution.
3. Write down solution
4. Advisor makes annoying and mystifying edit requests on irrelevant introductory stuff, while throwing out long complicated proofs (or experiments) student has spent many hours sweating over.
5. Make final edits and submit paper.

Most student figure out how to do step 2 , and eventually step 3. Step 5 is probably the first thing students learn how to do: Fix typos, edit latex, and generally do yak-shaving. But step 4 is perhaps the most mysterious part of the writing process for a new researcher, and the least structured. I call it "closing the deal" and it's really about going from a bag of results to an actual submittable paper. Let me elaborate.

1. Coming up with a problem. Of course coming up with a problem is the essence of the research process ("It's about the questions, not the answers", he shrieks). This takes experience and vision, and can often be changed by things you do in stage 4. I'll say no more about it here.

2. Solving a problem. This is the stage that everyone knows about. That's what we do, after all—solve problems ! This is where we drink lots of coffee, live "Eye of the Tiger" montages, get inspiration in our sleep, and so on. It often happens you don't exactly solve the problem you set out to attack, but you make many dents in it, solving special cases and variants. It's important to be flexible here, instead of banging your head against a wall head-on. At any rate, you either exit this stage of the project completely stuck, with a complete solution, or with a collection of results, ideas, and conjectures.

3. Writing it all down. Again, I could spend hours talking about this, and many people better than I have. It's a skill to learn in and of itself, and depends tremendously in the community you're in.

4. Closing, or getting to a submission. This is the part that's often the most critical, and the least understood—getting from 80 to 100 percent of a submission—it requires a different kind of skill. The overarching message is this: A paper tells a story, and you have to shape your results—their ordering, presentation, and even what you keep and what you leave out—in order to tell a consistent and clear story. (Before people start howling, I'm not talking about leaving out results that contradict the story; that would be dishonest. I'm talking about selecting which story to tell.) So you have a bag of results centering around a problem you're trying to solve. If the story that emerges is: "Here's a problem that's been open for 20 years and we solved it," then your story is relatively easy to tell. All you have to do is explain how, and using what tools.

But in general, life isn't that easy. Your results probably give some insights into the core of the problem: What parts are trivial, what directions might be blocked off, and so on.

Now you need to find/discover the story of your paper. You can't do this too early in the research process: You need to explore the landscape of the problem and prove some results first. But you shouldn't wait too long either; this stage can take time, especially if the story changes. And the story will change. One way of thinking about what you need for a conference submission is a relatively tight, compelling and interesting story. While the loose ends and unexplored directions are probably the thing most interesting to you and your research, they are best left to a conclusions section rather than the main body. What the body should contain is a well-thought out march through what you have discovered and what it says about the problem you're solving. In doing so, you will find yourself making decisions about what to keep, and what to leave out, and how to order what you keep.

And so, speculations need to be made into concrete claims or triaged. Experiments need to be ran until they tell a definite story. Introductions need to be made coherent with the rest of the paper. There's also an element of bolt-tightening. And all of this has to be done to serve the overarching story that will make the most compelling paper possible. The story can change as new results come in, or expand, or sometimes even die, which is rare. But there is a constant drumbeat of "Am I getting closer to a submission with a nice story with each step".

Telling a good story is important. For someone to appreciate your paper, cite it, or even talk about it (whether it's accepted, or on the arxiv) they have to be willing to read it and retain its results. And they'll be able to do that if it tells a clear story, which is not just a union of results.

Suresh Venkatasubramanian is an associate professor in the School of Computing at the University of Utah. He is currently a visiting scientist at the Simons Institute for Theoretical Computer Science and Google Inc. He spends his days plotting the takeover of the world by algorithms, especially geometric algorithms for large data problems.

Quantified Performance: Assessing runners with sensors

A look at how athletic performance can be measured outside of the laboratory.

By Christina Strohrmann and Gerhard Tröster

DOI: 10.1145/2541649

Sports scientists, trainers, and athletes often want to gain more insight into physical action and movements: Why can one athlete jump higher or run faster than the other? What are the major reasons for sustaining an injury? Is physical therapy actually helping to improve movement? Answering such questions can help sports professionals in their work and also allow for optimization of human performance. Yet current approaches to performance measurement rely on self-reports, visual observations, and video recordings for the analysis of an athlete's movement. However a more sophisticated approach exists. Using optical motion capture systems, where reflective markers are attached along the limbs, numerous infrared cameras capture the 3-D position of the markers. Computer models can then reconstruct a person's movements using a stick figure, making it possible to identify joint angles, angular velocities, and translational motions from the collected data.

While these measurement systems are highly accurate (they are capable of tracking movements of only a millimeter), they are not accessible to everyone. Moreover, they need a specific measurement environment and thus do not allow for unconstrained monitoring in the field. The high setup time also does not allow for assessments on a regular basis, and the large quantity of data to be processed does not allow for continuous long-term monitoring. These problems present a challenge to the measurement of sporting performance in everyday settings.

WEARABLES FOR MOVEMENT PERFORMANCE ANALYSIS

Wearable technologies offer a way to overcome the main drawbacks of current performance measurement systems. In particular, small sensors—most commonly accelerometers and gyroscopes—can be attached to the body to capture movement. These sensors can be used in everyday surroundings, are small and unobtrusive, and are potentially accessible to everyone due to their low cost and increasing availability. Many smartphones contain accelerometers and gyroscopes, making phones the most popular representation of a wearable computer to date.

In our research, we have been exploring the use of wearable technologies to analyze the performance of runners. We chose to monitor runners due to the increasing popularity of long-distance running (e.g. in half and full marathons) and the nature of the sport: Everyone can do it everywhere. Additionally, while many people opt for running as a way of keeping fit, most runners do not have a way of analyzing their movements to improve performance and reduce the risk of injury. This is potentially problematic because run-

Photograph by Steve Garner

ning carries a high injury risk, with 74 percent of runners suffering an injury in an average year [1]. Besides improper performance of movements, one contributing factor to injury was found to be fatigue [2]. When runners perform compensatory movements when fatigued, they increase the risk of sustaining an injury. Although performance monitoring could be used to minimize the risk of injury, state-of-the-art motion capture systems do not allow runners to be monitored while out running in the streets. Therefore, analyses are typically performed on treadmills and the results are then generalized to overground running. However, treadmill running does not adequately represent overground running [3]. Additionally, optical motion capture does not allow runners to be monitored during a prolonged run and cannot provide real-time feedback to runners. There is, thus, a clear opportunity to provide performance monitoring of runners in the field using wearable technologies.

ANALYSIS OF MOVEMENT PERFORMANCE USING ON-BODY SENSORS

Monitoring runners in the field requires the assessment system to have a number of characteristics. First, the system needs to be unobtrusive: The athlete should not be influenced in their movements, and they should be able to perform each movement as they would normally. The system also needs to be unconstrained and should not be reliant on a specific environment, allowing athletes to be monitored in their natural surroundings to better reflect their actual movement. The system should also be highly accurate, and the measurement accuracy should offer a sufficient signal to noise ratio to allow for detection of even small differences in performed movement. Finally, the system must be able to perform monitoring over the longer term, e.g. by capturing a full rehabilitation session or a full workout session.

To fulfill these requirements, we used the ETH Orientation Sensor (ETHOS) that was developed in previous work [4]. ETHOS is an inertial measurement unit (IMU) optimized for long-term recordings in the field. Each unit features a 3-D accelerometer, a 3-D gyroscope, and a 3-D magnetic field sensor. We developed two housing units (a flat housing and a bracelet) to allow for optimal attachment, as depicted in Figure 1. The round housing unit weighs 27 grams, and the flat housing unit weighs 22 grams including the sensor, the battery, and a switch. To test the efficacy of wearables for performance monitoring, we conducted experiments where we aimed to assess runners' skill level, technique, and level of fatigue [5, 6].

To explore differences in skill level, our experiments included runners of different experience, ranging from beginners to experts. In total 23 runners participated. To ensure we measured differences associated with skill level and not different velocities, we had all runners run on a treadmill with a velocity of 11.4 km/h. This allowed us to explore the parameters by which novice and experienced runners could be distinguished, as well as the number of sensors that would be required to make such a distinction in everyday settings.

With running being such a high injury risk sport and fatigue being a main contributor to injury, the later stages of our experiments used a protocol that ensured runners would experience fatigue. First, runners participated in an all-out test to assess their maximum aerobic velocity. Runners were then advised to run at 80-85 percent of this maximum velocity during a 45-minute run, and we explored how sensor data could be used to identify patterns of movement associated with fatigue. The experiment also allowed us to investigate differences in running technique; in this case, we explored foot striking patterns and how different patterns can be assessed using wearable sensors attached to the foot.

While many people opt for running as a way of keeping fit, most runners do not have a way of analyzing their movements to improve performance and reduce the risk of injury.

Each runner was equipped with 12 ETHOS units to monitor full-body movement. The measurement setups are depicted in Figure 1. We used questionnaires to see how runners felt while running with the ETHOS sensors. The questionnaires revealed runners did not feel restricted in their movements, suggesting to us the sensors and attachments were sufficiently comfortable and did not intrude on the runner's movements.

Figure 1: Runners were equipped with 12 ETHOS sensors.

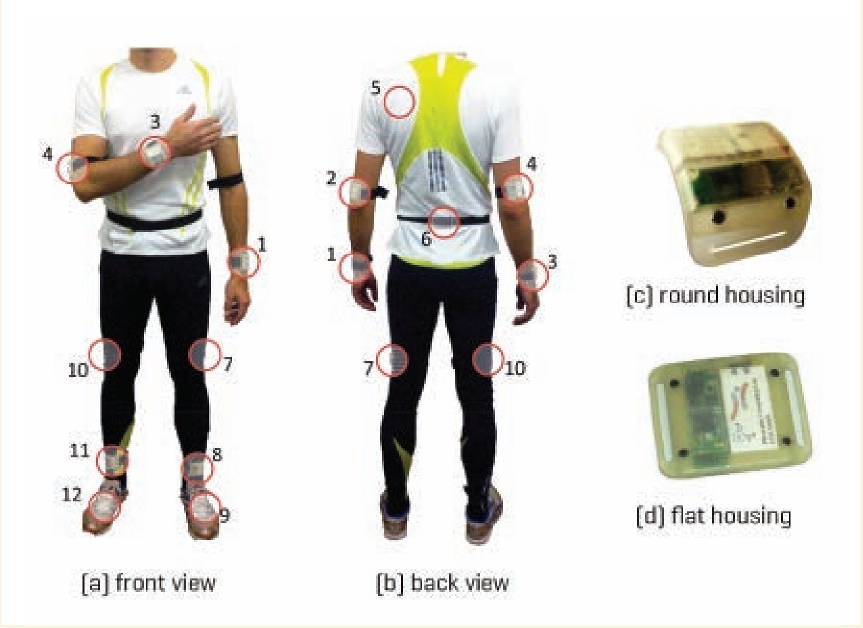

ASSESSMENT OF SKILL LEVEL

We found two parameters sufficed to distinguish between experienced and inexperienced runners: The amount of vertical oscillation, which refers to the amount of up and down movement measured at the hip, and foot contact duration. The latter is the time the foot remains on the ground during one gait cycle normalized by the gait cycle duration. In order to run fast, runners train to shorten foot contact duration, since more time in the air means longer flight time.

We calculated the vertical oscillation by a double integration of the acceleration and then subtracted the minimum from the maximum vertical exertion. (Note double integration introduces a quadratic error due to integration of noise. Therefore, the signal was reset at every detected foot strike.) The foot contact duration was calculated by first detecting foot strikes, indicated by a sharp peak in the foot's acceleration. Afterwards, acceleration remains close to 1g (Earth's gravity), meaning the foot is not moving. When the foot is being moved again, dynamic accelerations are superposed to the static acceleration of Earth's gravity. This means we can use a simple threshold algorithm to detect the end of foot contact. Foot contact duration was then normalized by step duration. We were able to show just two sensors are needed for the assessment of skill level: One on the hip and one on the foot. The results are depicted in detail in Figure 2a.

ASSESSMENT OF TECHNIQUE

Most runners are heel strikers, meaning they touch the ground with their heel first. While this is a very natural form of running, it slows runners down since every heel strike breaks a little. In addition, heel striking produces high impact on the knee. Therefore, most long distance runners train to use a midfoot striking pattern in which the

Figure 2: We assessed skill level (a), technique (b), and fatigue (c) from the data collected from runners using wearable technology. For skill level, we were able to show two sensors (on the foot and hip) suffice to distinguish between experienced and inexperienced runners. One sensor on the foot suffices to classify a runner's foot strike: heel, midfoot, or toe striking, depending on the direction of the foot's rolling motion. For fatigue we found runners of all skill levels didn't lift their heel as high when fatigued. We also observed changes among individual runners, e.g. dropping the shoulders with fatigue (lower right in the figure), which might indicate weak back muscles.

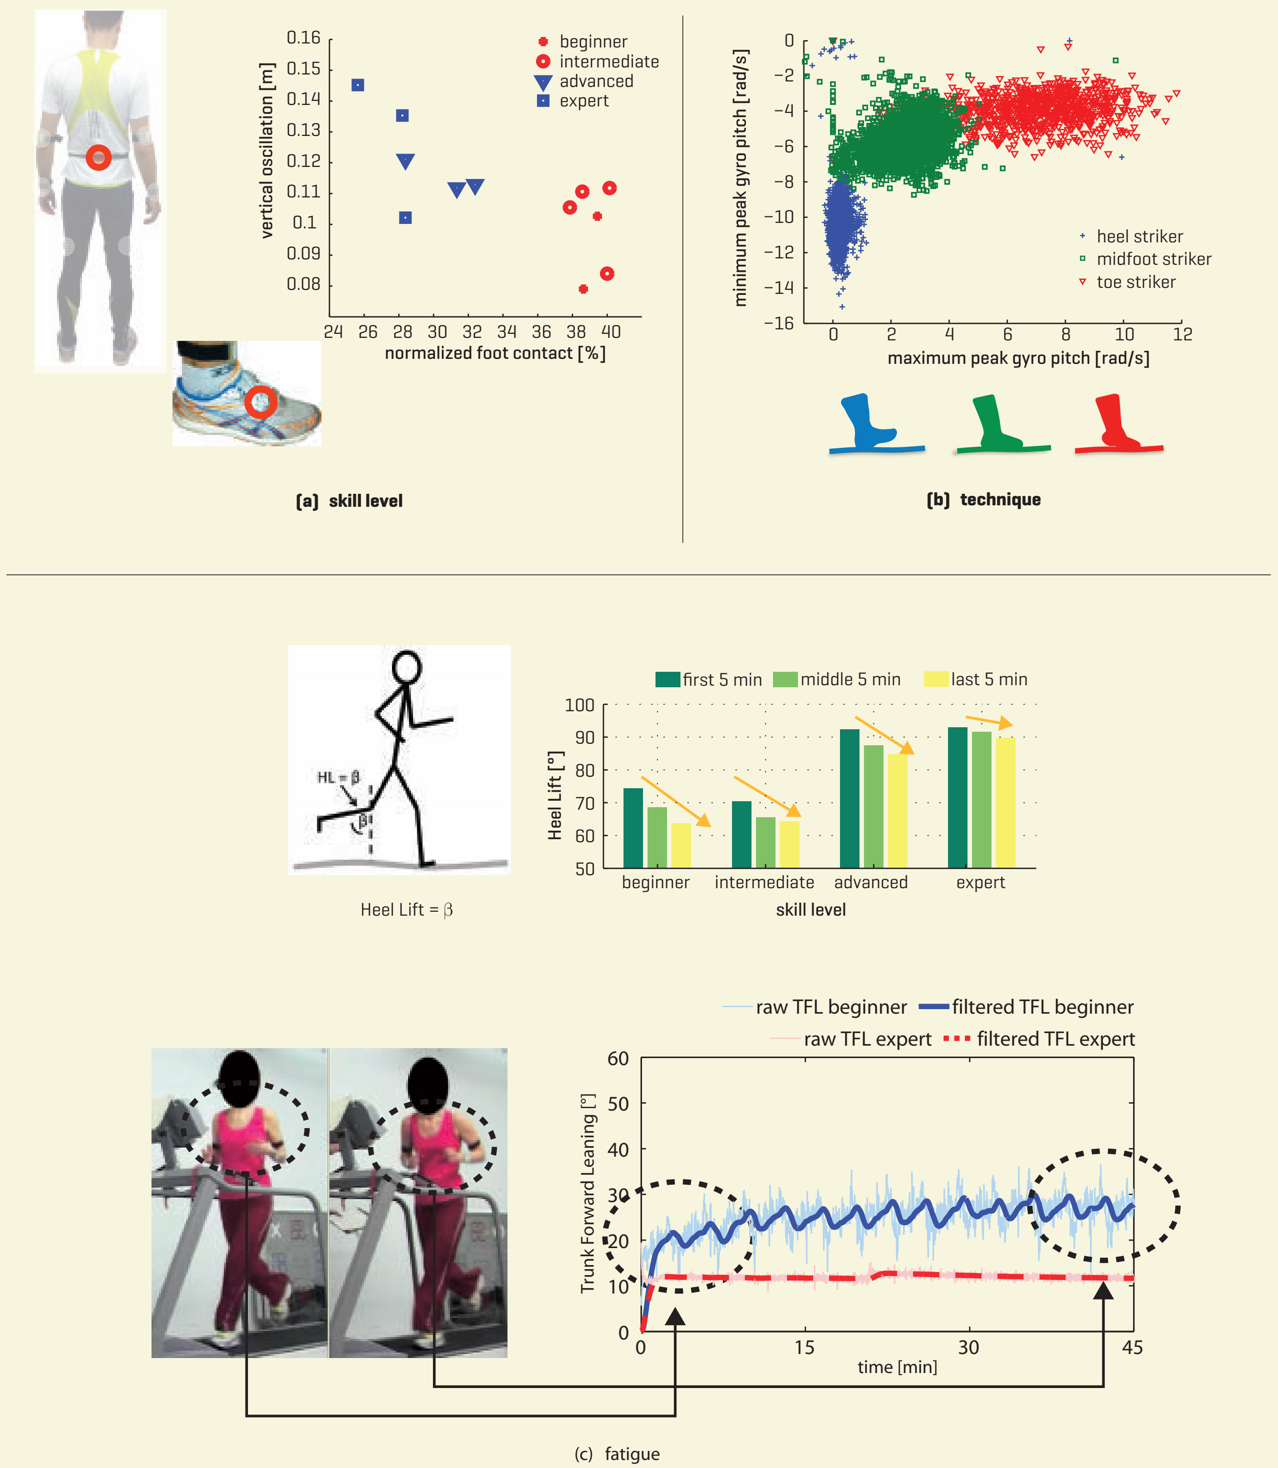

Figure 3: Feedback was provided to improve arm carriage while running. Running books advise to not let the arms cross the symmetry line as excess rotation increases the amount of wasted energy and stress on the lower back (top left figure). We implemented an online detection of false arm carriage on a smartphone. The smartphone vibrated when the faulty arm carriage was detected (right figure). In a user study we showed runners improved their arm carriage using the app. This improvement was comparable to that elicited by a verbal instruction.

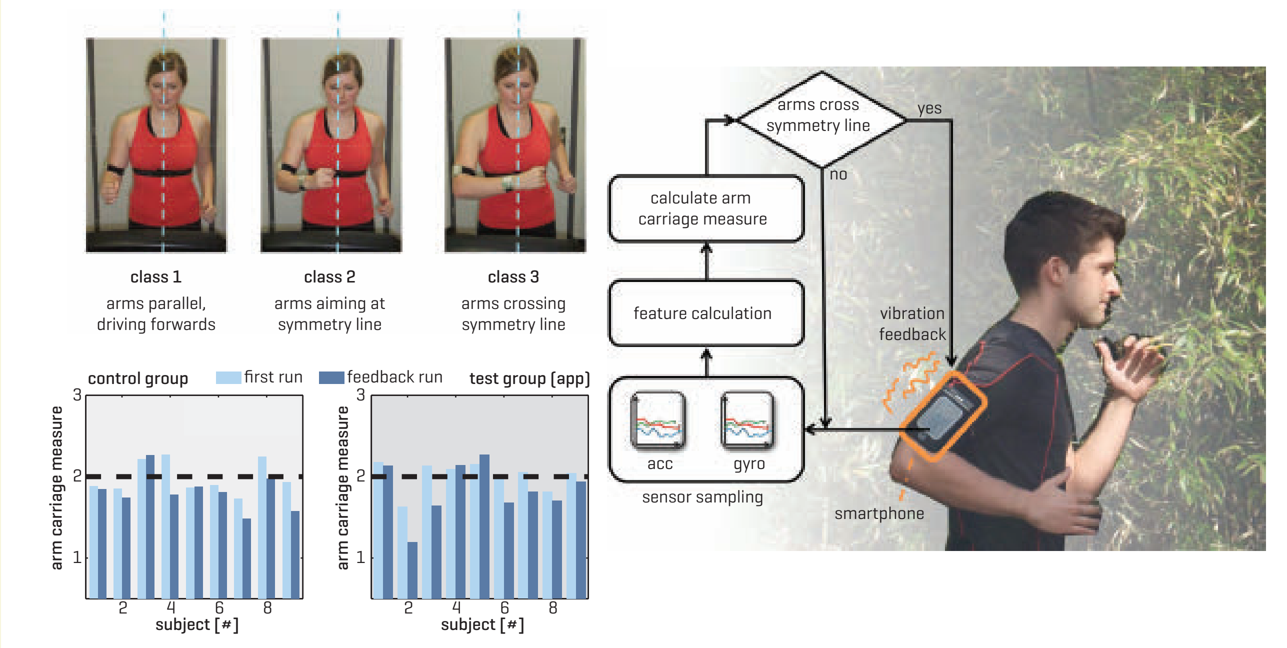

whole of the foot strikes the ground. This decreases the amount of impact on the knee, but slightly increases the stress on the calves and Achilles tendon. The third striking pattern, forefoot running, allows for fast running and is often observed during sprinting. However, it stresses the calves and Achilles tendon, increasing the risk of sustaining an injury.

The foot strike pattern can be assessed using a single sensor on a runner's foot. This was based on a simple analysis of the direction of rotation during foot strike: A heel striker rotates from his heel to his toes. A forefoot striker rotates the other way, from toes to heel, yielding an opposite sign in the rate of turn signal. A midfoot striker does not rotate much during foot strike. One sensor on the foot suffices to classify a runner's foot strike: heel, midfoot, or toe striking, depending on the direction of the foot's rolling motion. Data are presented in Figure 2b.

MONITORING OF FATIGUE

To monitor movement changes with fatigue, we calculated 10 established kinematic parameters (such as step frequency) from the sensor data and identified parameters that changed with fatigue for all runners, parameters that change for runners of distinct skill levels, and parameters that are dependent on an individual's running technique. As depicted in the upper area of Figure 2c, we observed the following; as they became fatigued, runners tended to lessen the lift in their heel. This was true for all runners irrespective of skill level. Additionally, we were able to identify individual differences from the sensor data and confirm these differences with observations from the videos. For example, one runner dropped her shoulders with fatigue, and we could also see her upper body wasn't as upright as the expert runner's (see Figure 2c, lower area).

Since running form and changes in fatigue are highly dependent on the individual runner, wearables offer a way to personalize running analysis in an everyday scenario.

Since running form and changes in fatigue are highly dependent on the individual runner, wearables offer a way to personalize running analysis in an everyday scenario. Runners would be able to not only track their progress in terms of running form, but also analyze their individual fatigue pattern. This would then allow runners and trainers to identify possible interventions for injury prevention and the improvement of performance. For the example of shoulder dropping when fatigued, a runner might opt for strength training of the back in addition to their ordinary running exercises.

FEEDBACK PROVISION

In further work, we investigated how wearables could provide real-time feedback to runners. A common mistake while running is to perform a large shoulder rotation, which wastes energy and increases the strain on the lower back. Trainers typically advise runners to move their arms in parallel with the direction of walking without crossing the symmetry line, as depicted in Figure 3.

In a study with 10 participants, we investigated different sensor positions on the arm and the back for the automatic detection of arms crossing the symmetry line. Feedback could then be provided to help the runner improve his or her form. We found the best sensor position was on the upper arm. Since the upper arm is a common location for music players and smartphones while running, we investigated the use of a smartphone for arm carriage monitoring and feedback provision.

The feedback was implemented as an Android application, using the phone's integrated accelerometer and gyroscope to monitor arm carriage. We created a detection algorithm and trained it using data from the participants in our study. The algorithm detects the amount of rotation along the vertical axis of the arm and the amount of elevation of the elbow, as measured by the acceleration sensor. When the algorithm detects that the arms have crossed the symmetry line, the phone offers feedback to the runner in the form of vibration. A flowchart of the app is shown in Figure 3.

To evaluate our application, we performed a user study with 20 participants. Each runner performed two 20-minute runs. The first was a control run and the second was an experimental run in which runners were randomly assigned to one of two feedback groups: app feedback (test) or verbal feedback (control). The app group received vibration feedback from the smartphone, and the verbal feedback group received verbal prompts from a human trainer on how to perform correct arm carriage. In the latter case, the phone was used to measure the movement and did not provide any feedback. Figure 3 depicts our results. In short, we found runners improved their arm carriage in both groups.

While using a smartphone with force feedback might not yield significantly better results over verbal instruction from a trainer, a smartphone is available to everyone. Our system offers a convenient way of monitoring performance for those who do not have personal trainers. This is especially helpful since runners could use the force feedback method to train on their own. A questionnaire also revealed runners liked the system and many stated they would be interested in using the tool on a regular basis.

CONCLUSION

Movement analysis provides a valuable tool for performance assessment in running. To date, the state-of-the-art approach uses optical motion capture systems and ground reaction force measurements for movement assessment. While these provide high accuracy, they are restricted to an instrumented environment and need high setup times. With our findings, running analyses can be made available for all runners, especially if they don't have access to a personal trainer. Additionally, the use of wearables can help to further understand the complex relationship between running technique, injury, and running economy, a topic that is still not very well understood [7].

This article focuses on the use of wearable technologies for performance monitoring in running, yet their application is not restricted to one type of exercise. Wearables have already been used in other sports such as swimming or snowboarding [8, 9]. Each application scenario presents different challenges. For swimming, wearables must be not only sweat resistant but also entirely waterproof, which is a major research challenge. Conversely, snowboarding requires a feedback method that does not distract the rider from the slope and from other riders nearby. Another major research area is the application of wearable technologies in the healthcare domain, especially for rehabilitation procedures that tend to rely on subjective judgements made by doctors (see page 33). It would be useful to have sensors that could be worn at home to monitor rehabilitation progress, e.g. by tracking movements of the body from day-to-day to show how the movement of limbs has improved over time. Sports and rehabilitation are just two of the ways in which wearables could be used to monitor movements and improve physical wellbeing—we look forward to seeing what others achieve in the future.

References

[1] Daoud, A.I. , Geissler, G.J., Wang, F., Saretsky, J., Daoud, Y.A., and Lieberman, D.E. Foot strike and injury rates in endurance runners: A retrospective study. *Medicine and Science in Sports and Exercise* 44, 7, (2012), 1325–34.

[2] Mizrahi, J., Verbitsky, O., Isakov, E., and Daily, D. Effect of fatigue on leg kinematics and impact acceleration in long distance running. *Human Movement Science* 19, 2, (2000), 139 – 151.

[3] Riley, P.O. et al. A kinematics and kinetic comparison of overground and treadmill running. *Medicine and Science in Sports and Exercise* 40, 6 (2008), 1093-1100.

[4] Harms, H., Amft, O., Winkler, R., Schumm, J., Kusserow, M., and Tröster, G. ETHOS: Miniature orientation sensor for wearable human motion analysis. In the *Electronic Proceedings of the Ninth IEEE Sensors Conference, IEEE Sensors* (Waikoloa Village, Hawaii, Nov. 1-4). IEEE, New York, 2010, 1037–1042.

[5] Strohrmann,C., Harms, H., and Tröster, G. What do sensors know about your running performance? In *Proceedings of the 15th IEEE International Symposium on Wearable Computers* (San Francisco, CA, June 12-15). IEEE, New York, 2011, 101-104.

[6] Strohrmann,C., Harms, H., and Tröster, G., Hensler, S., and Müller, R. Out of the lab and into the woods: Kinematic analysis in running using wearable sensors. In *Proceedings of the 13th ACM International Conference on Ubiquitous Computing* (Beijing, Sept. 17 – 31). ACM Press, New York, 2011, 119-122.

[7] Burgess, T.L. and Lambert, M.I. The effects of training, muscle damage and fatigue on running economy: review article. *International SportMed Journal* 12, 1 (2011), 363–79.

[8] Bächlin, M., Förster, K., and Tröster, G. SwimMaster: A wearable assistant for swimmer. In *Proceedings of the 11th International Conference on Ubiquitous Computing* (Orlando, FL, Sept. 30 – Oct. 3). ACM Press, New York, 2009, 215–224.

[9] Spelmezan, D., Jacobs, M., Hilgers, A., and Borchers, J. Tactile motion instructions for physical activities. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems, CHI 2009* (Boston, April 4-9). ACM Press, New York, 2011, 2243–52.

Biographies

Christina Strohrmann received the Dipl.-Ing. degree in information technology from the University of Kaiserslautern, Germany, in 2010. She joined the Wearable Computing Group at ETH Zurich as a reesearch and teaching assistant in 2010. Her research interests include movement analysis in sports and rehabilitation using small body-worn sensors.

Gerhard Tröster received the Dipl.-Ing. degree in electrical engineering from Darmstadt and Karlsruhe, in 1979, and the Dr.-Ing. degree from the Technical University Darmstadt, Darmstadt, Germany, in 1984. He was involved in the research on design methods of analog/digital systems in CMOS and BiCMOS technology for eight years at Telefunken (atmel), Heilbronn. Since 1993, he has been a full professor of electronics at ETH Zurich, heading the Electronics Laboratory. At ETH, he established the multichip module (MCM) electronic packaging group. In 2000, he founded the Wearable Computing Laboratory, ETH, where he was involved in interdisciplinary research combining IT, signal processing, electronic platforms, wireless sensor networks, smart textiles, and human-computer interaction.

Fitness Trackers

Digital activity sensors are no longer confined to research labs; they're in the wild and they come in lime green. They offer the promise to improve our health and even to affect the ways that we interact with others.

By Andrew Miller

DOI: 10.1145/2543611

Wireless fitness trackers are ready for their closeup. Today, you can buy a digital pedometer from companies like Jawbone, Nike, Withings, and Fitbit, with new models released seemingly every week. There are even crowd-funded trackers, like the Misfit Shine. Smartphone-based apps can now estimate daily movement, expanding availability even further. This new generation of fitness sensors is robust, colorful, and networked. Friends can cheer you on when you take your Nike Fuelband for a run; you can track your minute-by-minute data with the FitBit. Pervasive fitness tracking has truly "exited the cleanroom" and entered the wild [1].

As a research community, human-computer interaction (HCI) has been preparing for this day. We've been studying how people interact with on-body sensors for almost a decade now, using custom prototypes, expensive niche devices, and early digital pedometers. But these all required a certain degree of handholding from the research team, and the turnaround time from sensing to reporting back was hardly instantaneous. As in any human-centered investigation of the future, the present kept getting in the way—people just weren't used to wearing computers in their pockets, much less sharing their daily lives online. The "qunatified self" movement—in which people self-monitor as much as they can of their daily habits, moods, exertion, and food—took to fitness trackers early on, testing different form factors and data presentation styles and pushing the limits of the technology as only early adopters can. However, as early adopters, they can tell us only so much about how technologies will be used by the population at large.

Now, all the pieces are ready. Smartphone adoption is broad, and fast becoming near-universal; Facebook has a billion users; and you can buy a networked wireless pedometer with a three-month battery life that will survive a trip through the washing machine. Oh, and it comes in lime green.

This is an exciting time, because it calls for a different set of research skills. Early-stage researchers focused on how to make the technology work, but today's pervasive fitness researchers can also focus on the "why." We can now run longer and larger studies in more diverse surroundings, and we can focus on new problem domains beyond the individual. We can work with health promotion researchers to test the health impact of our systems, and we can work with communities to understand the social and cultural impact of introducing on-body sensors into everyday life.

We can also shift our research questions "up the stack." That is, we're now able to study the fitness tracker as more than just a personal data-gathering tool; we can now treat it as a social and cultural artifact. We can embed trackers into new kinds of socio-technical systems, ones far different from those we'd construct for lab studies or feasibility deployments. And we can begin to study how fitness trackers might help people manage their everyday health as individuals and communities.

We must adopt different methods, taking a more "in-the-wild" and human-centered approach to our research.

FIGHTING CHILDHOOD OBESITY

In my research, for example, I work with middle school students (11 to 14 years old) on technologies for obesity prevention. These kids present some unique challenges and opportunities for fit-

Photograph by Brandon Shigeta

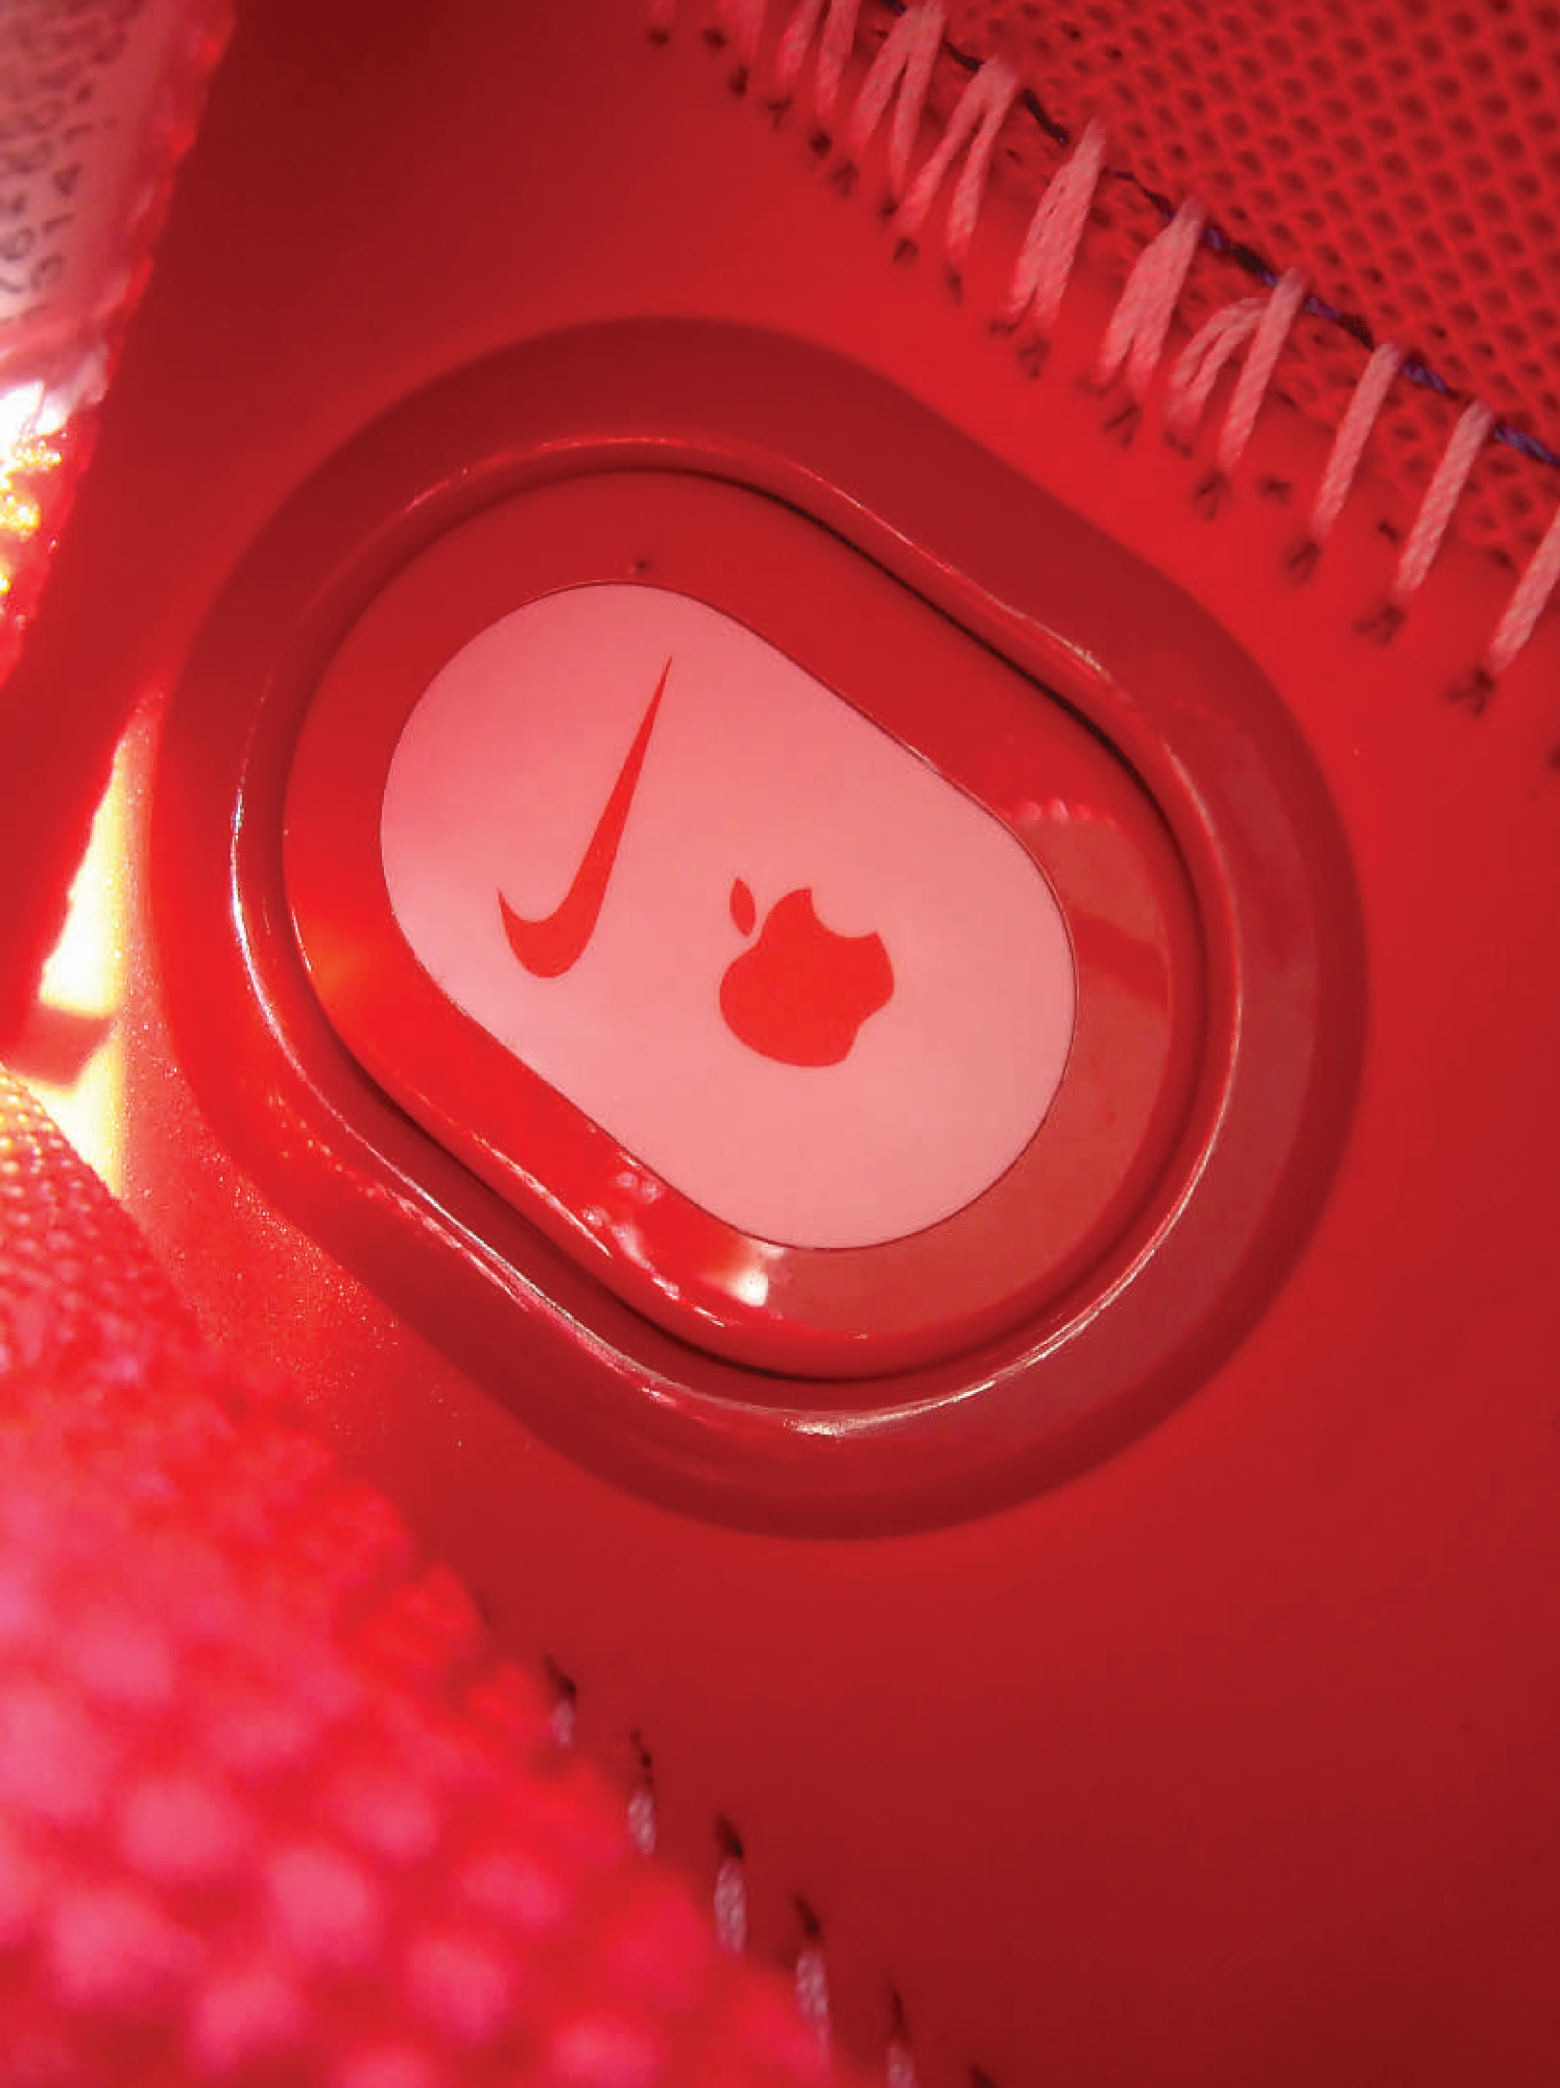

ness tracker studies. Like all children their age, they want simultaneously to stand out and to fit in. Their identities and preferences are tightly bound to what their peers say and do. The boys and girls I work with have starkly different attitudes toward competition. And nothing is more boring to these kids than a bar chart of their own physical activity. They also live in poor, urban neighborhoods, where walking outside can be dangerous. Until this year, few of them had smartphones, and demographically they're in real danger of becoming overweight. About 60 percent of adults in their community are overweight or obese. Anecdotally, they appear to be unsupervised after school; many live in single-parent households and several have fathers in prison.

You might think it crazy to conduct a study in such a swirl of social, cultural, and economic factors, which could overwhelm or derail a fitness tracker study. Instead they become features of the design space. I knew my system had to be social: It had to provide a way for kids to motivate each other while being minimally demotivating to less-active kids. It had to be school-based: Creating an after-school program and working with administrators and teachers enabled me to work with a group of kids who saw each other daily, making it possible to get them together for weekly deployment meetings. It had to be designed with the their help: I'm not their age and I don't live in their community, so I involved kids from the school throughout the design process. Finally, the chosen fitness tracker had to be robust and kid-friendly. Fortunately, a few months before my deployment Fitbit released the Fitbit Zip, so I was able to focus on the social and behavioral effects and leave the hardware hacking to others.

Even crazier, it appears to have worked. The system I created, StepStream, pulled students' individual daily step-counts into a social network site. Kids earned activity points they could spend on a social game, and they met weekly to chat on the site and play the game. They also had access to the site between meetings and wore their pedometers throughout the month-long deployment.

This study truly was "in the wild." Kids took their pedometers everywhere; a quarter of them even wore the pedometers to bed despite the lack of sleep tracking in the pedometers themselves. (When was the last time someone wore your research project to bed?) I'm also learning interesting things about the interaction of the social features, individual motivation, and online/offline interactions that would not have shown up in a more controlled setting.

However, the fact that these wireless fitness trackers have been commercialized doesn't mean the supporting infrastructure has disappeared. During my most recent deployment, I had to drive to the school three times to reset the base station after a power outage. Of the 42 pedometers I handed out, I replaced or repaired more than 20. And the server I was using to host the system suffered a 19-hour outage (fortunately, mostly overnight). When it restarted, the server was in a different time zone, forcing me to adjust the timestamps for 14 days of system logs.

As far as health outcomes, we still have a way to go. Forty participants in a month-long study seems long to HCI researchers, but to the public health community it's a small pilot study. Changes in physical activity behavior may take months or years to stabilize, and require deep psychological changes. Participants have to change their identities to see themselves as healthier, more active people. The gold standard in health research—the randomized controlled trial (RCT)—demands a longer intervention and more technological stability than most HCI studies can promise.

But we are making great progress from a computing research standpoint. Ubiquitous computing theory from a decade ago is finally impacting people in daily life. For example, in his 2001 book *Where the Action Is*, Paul Dourish made the case that tangible and social computing were on a collision course, and were actually two sides of the same coin: embodiment. Embodied technologies, per Dourish, are situated in the physical and social world that we inhabit, and the more embodied they are the more of our context they share. Fitness trackers are embodied interaction made real. In my research, I situated pedometers within a social context (an urban middle school) and a technical system (a social website), but that's just the start.

For example, the fitness trackers, themselves, could facilitate social interaction more directly or become social actors in themselves. The Fitbits used in my study fed information in one direction: into the cloud. The devices couldn't react if two participants were near each other, or behave differently when placed on the hip versus the chest. They did display Tamagotchi-style faces in reaction to recent activity levels, but these proved inscrutable to the kids and, in any case, only reflected individual activities.

Today's fitness tracker research will also help us prepare for the next wave of wearables and on-body networks. The hardware hackers haven't stopped; they've just moved to more exotic tech like all-day heart rate monitors and stickers that sense your blood pressure—all communicating to each other and to a remote activity profile in the cloud. But until that day comes, there's plenty to be done now.

Fitness tracker research is at a crossroads. As computing researchers, we can now study how these technologies will be used in the daily lives of millions, and our research has the potential for meaningful impact on important societal issues. To push the state of the art forward, we must adopt different methods, taking a more "in-the-wild" and human-centered approach to our research. This research may offer fewer clean proscriptions, but its rich descriptions of technology in use will position us well for the next phase: seeking out collaborations beyond computing. Experts from domains such as healthcare and education know how to show efficacy, but will need our guidance to understand the role of technology as it reshapes their research as well.

It's a comforting thought: We're at the crossroads, but we're not alone, and we don't have to do it all. We can ask new kinds of questions and work with new collaborators, one step at a time.

Reference

[1] Carter, S. et al. Exiting the cleanroom: On Ecological validity and ubiquitous computing. *Human-Computer Interaction* 23, 1 (2008), 47–99; doi:10.1080/07370020701851086.

Biography

Andrew Miller is a Ph.D. candidate in human-centered computing at Georgia Tech. He holds an M.S.-HCI (also from Georgia Tech) and a B.A. in cognitive science from Occidental College.

Tracking How We Read: Activity recognition for cognitive tasks

Using activity recognition for cognitive tasks can provide new insights about reading and learning habits.

By Kai Kunze

DOI: 10.1145/2538691

Traditionally, research in activity recognition has focused on identifying physical tasks being performed by the user using elaborate, dedicated sensor setups in the lab. In recent years, physical activity recognition has become relatively mainstream. As industry begins to apply advances in activity recognition research, we are seeing more and more commercial products that help people track their physical fitness—from simple step counting (e.g., Fitbit One, Nike Fuelband, Withings Pulse), to recording sports exercises (e.g., Runkeeper), to monitoring sleep.

While the problem of physical activity recognition has been well explored, the ability to detect cognitive activities is an open area with many challenges. This exciting new research field, cognitive "quantified self," is opening up new opportunities for graduate students at the intersection of wearable computing, machine learning, psychology, and cognitive science.

MOBILE COGNITIVE SENSING

An obvious way to track cognitive tasks is to directly analyze brain waves. In practice, however, this requires either bulky equipment (sometimes the size of a room) or quite expensive procedures (e.g., functional magnetic resonance imaging and electrocorticography). For directly sensing brain activity, electroencephalography (EEG) and functional near-infrared spectroscopy (fNIR) seem to be the most promising for mobile usage, as both are relatively low cost. Mobile, commercial EEG systems are already available or close to release, such as the Emotiv Insight Neuroheadset shown on the opposite page. Unfortunately, EEG and fNIR suffer from several problems, including poor spatial resolution, a large degree of noise, and costly inference. Also, good

Image courtesy of EmotivInsight.


EMOTIV
EMOTIV

Figure 1. Eye gaze analysis while reading. Left: A user with the SMI iView X eye tracker reads a text. Center: The user's eyes illuminated by infrared light as recorded by the eye tracker to record the eye gaze. Right: The user's eye gaze mapped onto the document she's reading. The lines represent saccades, and the circles are where users fixated their gaze. The circle size signifies the duration of the fixation.

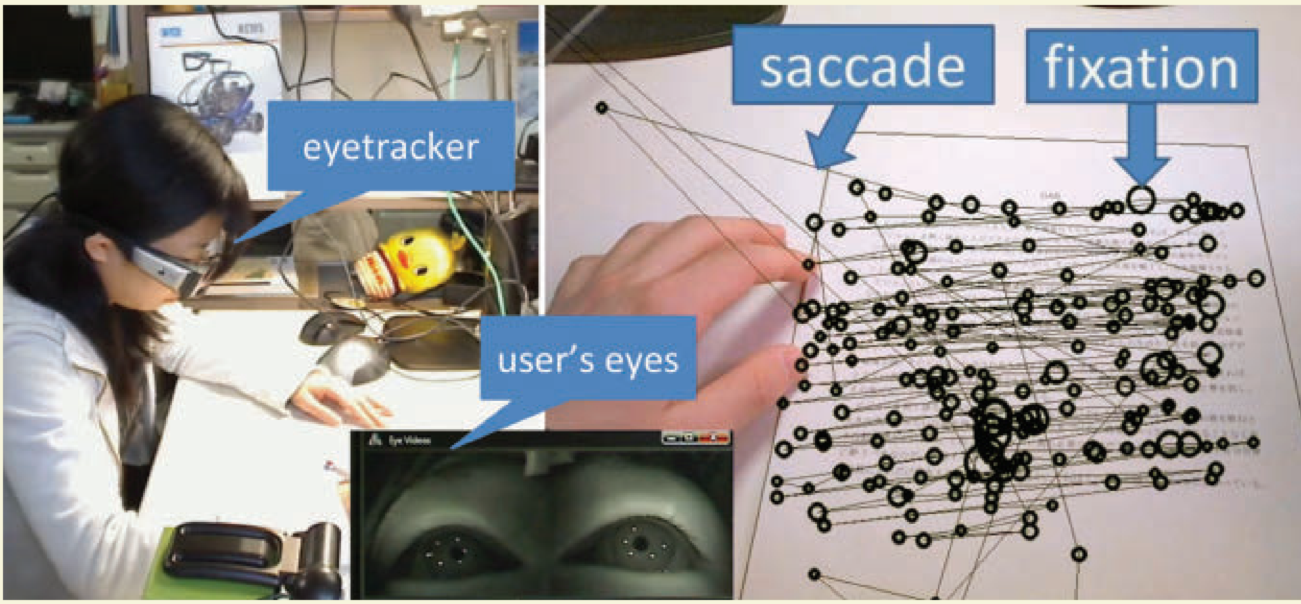

recognition results seem very sensor placement- and user-dependent.

An alternative is tracking eye movement, or "eye gaze," which has a strong correlation with cognitive activities. One challenge is eye gaze also correlates with the user's emotions and vitality, as well as environmental factors. Therefore, separating the cognitive tasks from the rest can be a nontrivial problem.

Two standard ways to implement mobile eye tracking are optical tracking and electrooculography (EOG). EOG measures the electronic potential in the eye between the cornea and retina (they can be represented as a dipole). If the eye moves, we see a change in electronic potential. Optical tracking systems usually use stereo cameras and infrared light, which is reflected by the eye. Computer vision techniques can be applied to estimate the gaze as a sequence of fixations (focus of the gaze on a single location at a time) and saccades (fast movements between fixations). Most systems use some feature-based computer vision models. Usually, the iris boundaries are modeled as circles in 3-D space. The normal vectors passing through the center of these circles estimates the gaze direction .

EOG is relatively low cost. However unlike optical tracking, EOG requires electrodes touching the skin near the eye. Additionally, with EOG, we can detect reading and other eye gestures in a mobile setting, but can't pinpoint where the user is looking [1]. Optical eye tracking can provide higher detail regarding gaze estimation and does not require direct skin contact, just more complex, expensive algorithms.

TRACKING READING HABITS

In our research, we mostly use optical eye tracking, with the occasional usage of first-person vision (a camera worn on the user's head) and EEG. Our main research focus is on tracking reading in mobile settings. By reading we mean the cognitive process of decoding letters, words, and sentences. We chose to study reading, a very interesting cognitive task, as it is a fundamental human technique used to acquire information. Furthermore, increased reading is positively correlated with improved general knowledge and language skill. Despite this, there are only a few existing studies addressing reading detection in real-life environments.

Our research involves quantifying various reading activities, like determining how much you read, what you read, and how much you understand.

Quantifying how much is read. First we implemented the Wordometer. Analogous to a pedometer counting the number of steps a user takes, the Wordometer estimates the words a user reads using the eye gaze recorded by a mobile eye tracker and document image retrieval [2].

We used eye gaze to recognize if a user is reading or not. While reading, we detected line breaks. Using these line breaks, we estimated the word count by simply multiplying the average number of words per line in the document by the detected line breaks (see Figure 1).

A more sophisticated algorithm uses support vector regression. We tested our algorithms on 10 users reading 10 documents and engaging in several not-reading activities. The simple version gives an average error rate of 13.5 percent. The more sophisticated word count algorithm reaches an average error rate of 8.2 percent (6.5 percent if one test subject with abnormal behavior is excluded). This is reasonably close to a pedometer error rate, which is between 3-10 percent. It's good enough to gain first insights into people's reading behavior.

Detecting what is read. We also investigated what you read using eye gaze. Tracking the document types enables us to gain more insights about the expertise level and potential knowledge of users—in the form of a reading log tracking and improved knowledge acquisition.

How often a user reads specific document types can provide insights into interests (e.g., comic versus belletristic) or language expertise and skills (e.g., computer vision textbooks versus English literature books). As a first step, we evaluated whether different document types can be automatically inferred from eye gaze [3].

We evaluated our approach using a combination of novel gaze features in a user study with eight participants and five Japanese document types: novel, manga, fashion magazine, newspaper, and textbook. We achieved a recognition performance of 74 percent using user-independent training. Figure 2 shows two of the document types with the corresponding characteristic eye gaze data.

Evaluating reading comprehension. An even more interesting question is whether we can estimate text comprehension and the level of expertise of the reader using eye movements. In an initial study, we focused on assessing second-language skills in students [4].

Wearing the mobile eye tracker, the participants read several text comprehension sections from a standardized English test, answered questions, and afterwards highlighted difficult words. Looking at the frequency of fixations, we can determine all difficult words marked by the user. We currently try to estimate the users' TOEIC score (a standardized English test for Japanese students) based on the results of the eye gaze. Perhaps in the future simply reading a passage of text might be enough to assess language skill, making standardized tests a thing of the past.

LOOKING AHEAD

Based on recently filed patent applications, Google, Apple, and other tech companies might already be experimenting with eye tracking prototypes. There are also several research groups

Figure 2. Detecting document types users are reading. Left: Scene images of the eye tracker while a participant was reading a textbook in the lecture hall (top), a fashion magazine at home (bottom). Right: Characteristic eye gaze patterns for both textbook (top) and fashion magazine (bottom).

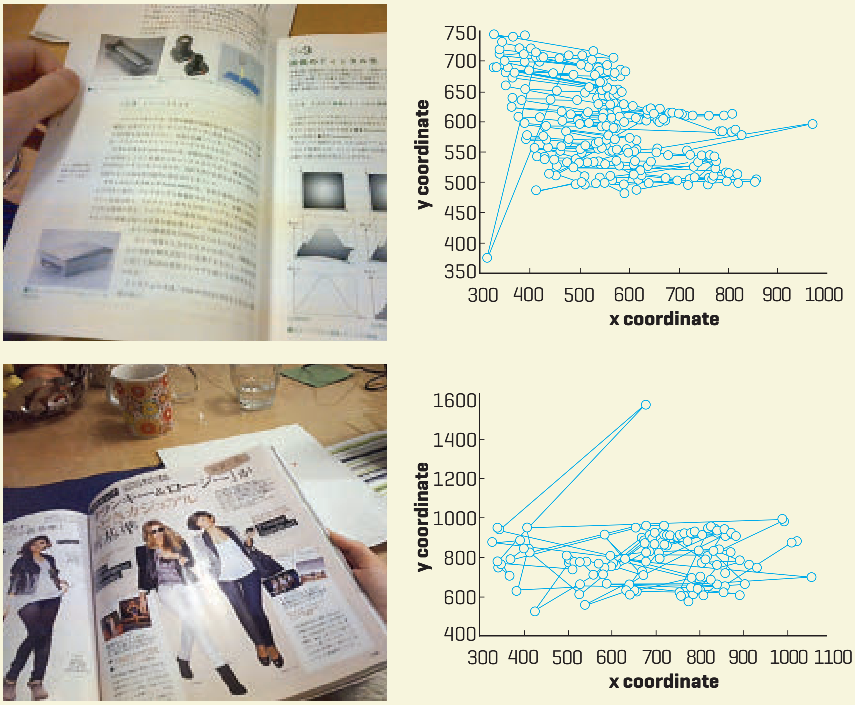

Figure 3. Interface mockup for a hypothetical Readit service that records and analyzes reading habits, similar to fitness services such as Fitbit.

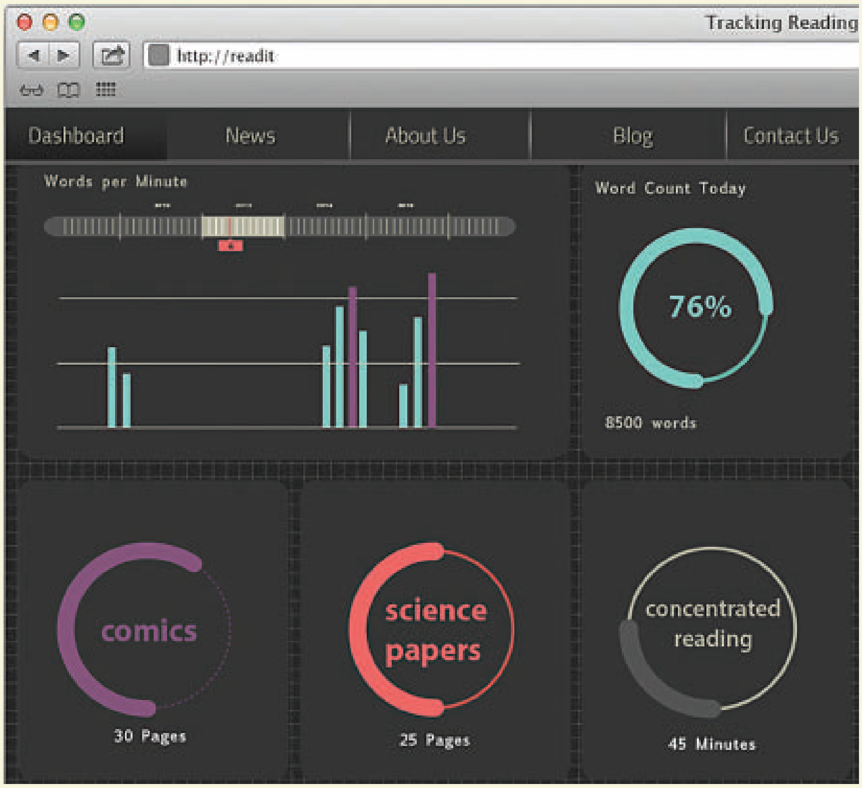

While the problem of physical activity recognition has been well explored, the ability to detect cognitive activities is an open area with many challenges.

working on low-cost, coarse-grain eye tracking using the front cameras on tablets and smartphones [5].

Although commercial eye trackers are currently still expensive, higher demand and increased availability of required hardware (cameras and infrared light) should drive prices down. Tech-savvy makers can already find build-it-yourself instructions for building cheap eye trackers on the Web [6]. As a first step, we can imagine a service similar to the physical activity tracking applications available today (e.g., Fitbit, Whitings, etc.), giving users a detailed overview about what they read and how much (see Figure 3).

Tracking reading habits can revolutionize education, as it enables students to evaluate their progress and receive personalized help for their experience level and particular needs. It will have a strong impact on publishing: Readers will be able to select documents according to their preferences (e.g., "Give me a book that lifts my mood"), and authors will be able to assess potential problems with their scripts (e.g., "Most users had trouble understanding this sentence").

Yet, there is also a dark side to this technology. Cognitive task recognition is not free of severe privacy concerns. It's a scary prospect that somebody could discern your expertise level on specific subjects or your feelings toward a topic. Therefore, users should be careful about granting applications access to the sensors on their mobile devices (e.g., the front-facing camera on your tablet), and researchers should be careful to identify potential misuse and privacy issues early on in order to communicate and ideally prevent them.

TOWARD A COGNITIVE "QUANTIFIED SELF"

This aritcle focuses on recognizing cognitive activities related to reading. However, this approach can be applied to any mental task [7]. For example, we can start to evaluate media intake and influence on knowledge acquisition tasks. We can assess when and why people lose interest in movies, TV shows, talks, or live performances, and see how to improve their narrative and user engagement. The same holds for software, both for recreation and business. Furthermore, we can find out about the impact of things like food choices, sleeping habits, and study breaks on cognitive performance. The possibilities are endless.

References

[1] A. Bulling, J., and Gellersen, H. Multi-modal recognition of reading activity in transit using body-worn sensors. *ACM Trans. Applied Perception 9*, 1 (2012).

[2] Kunze, K., et al. The Wordometer—Estimating the number of words read using document image retrieval and mobile eye tracking. In *Proceedings of the 12th Int'l Conf. Document Analysis and Recognition* (Washington, D.C., Aug. 25-28). IEEE, 2013.

[3] Kunze, K., et al. I know what you are reading—Recognition of document types using mobile eye tracking. In *Proceedings of the 17th Ann. Int'l Symp. Wearable Computers* (Zurich, Sept. 8-12). ACM Press, New York, 2013, 113-116.

[4] Kunze, K., et al. Towards inferring language expertise using eye tracking. *CHI '13 Extended Abstracts on Human Factors in Computing Systems* (Paris, Apr. 27-May 2). ACM Press, New York, 2013.

[5] Kunze, K., et al. My reading life: Towards utilizing eye tracking on unmodified tablets and phones. In *Proceedings of UbiComp '13 Adjunct* (Zurich, Sept. 8-12). ACM Press, New York, 2013, 283-286.

[6] Lukander, K., et al. OMG!: A new robust, wearable and affordable open source mobile gaze tracker. In *Proceedings of the 15th Int'l Conf. Human- Computer Interaction with Mobile Devices and Services* (Munich, Aug. 27-30). ACM Press, New York, 2013, 408-411.

[7] Kunze, K., et al. Activity recognition for the mind: Toward a cognitive quantified self. *IEEE Computer.* 2013 (to appear).

Biography

Kai Kunze is an assistant professor in the Department of Computer Science and Intelligent Systems at Osaka Prefecture University, Japan. He holds a Ph.D. from University Passau, Germany under the advisement of Prof. Lukowicz (DFKI, Kaiserslautern). For more information visit http://kaikunze.de.

Toward Smartphone Assisted Personal Rehabilitation Training

When utilizing internal sensors, modern smartphones are inexpensive and powerful wearable devices for sensor data acquisition, processing, and feedback in personal daily health applications.

By Gabriele Spina and Oliver Amft

DOI: 10.1145/2544048

For the first time in history, our generation and future generations will no longer be young. It is estimated that human life expectancy in the Stone Age was around 20-34 years. We can consider this as the natural life expectancy at birth for our species. However, nowadays, those born in Japan can expect to live 83 years. This implies there has been roughly a tripling of life expectancy for humans in the last few thousand years, which vhas dramatically altered the way societies and economies work.

Aging can be viewed as a triumph of development rather than any evolutionary changes in human biology: People are living longer thanks to technological and medical advances, better healthcare, education, and economic well-being. With increasing healthcare costs and a shortage of medical professionals, we are seeing a paradigm shift from hosting chronic patients in hospitals toward managing patients in their own home environment.

While deaths due to major diseases (such as AIDS/HIV and heart failure) are on the decline, the worldwide prevalence and related deaths of chronic diseases—such as chronic obstructive pulmonary disease (COPD), diabetes, and cardiovascular diseases (CVD)—are continually increasing (see Figures 1 and 2). COPD is predicted to become the third leading cause of mortality by 2030 [1]. It is estimated that 210 million people have COPD worldwide and 10.4 percent of the population older than 40 years have moderate to severe COPD that results in airflow limitation and significant extra pulmonary effects (e.g. muscle weakness and osteoporosis) [2]. Patients suffering from COPD have difficulty breathing and develop "air hunger." Breathlessness is a common occurrence forcing patients to avoid physical activities and enter into a vicious cycle: By exercising less, their muscles become weaker and less efficient; patients become more breathless and then gradually avoid exercising altogether.

How can COPD patients break this cycle and increase life expectancy?

Exercise training is a well-recognized method to treat symptomatic patients with COPD; physical activity programs appear essential to safely improve health state, including exercise capacity, functional status, health-related quality of life, peripheral muscle force, and physical activity in daily life. For example, generally healthy people can regularly jog and run, and even over-train, without immediate health consequences. In COPD patients, both over-training and undertraining can lead to the quick and detrimental worsening of health conditions, resulting in exacerbations, hospitalization, or death. For this reason,

Figure 1. Estimated mortality rates due to different diseases (World Health Organization July 2013).

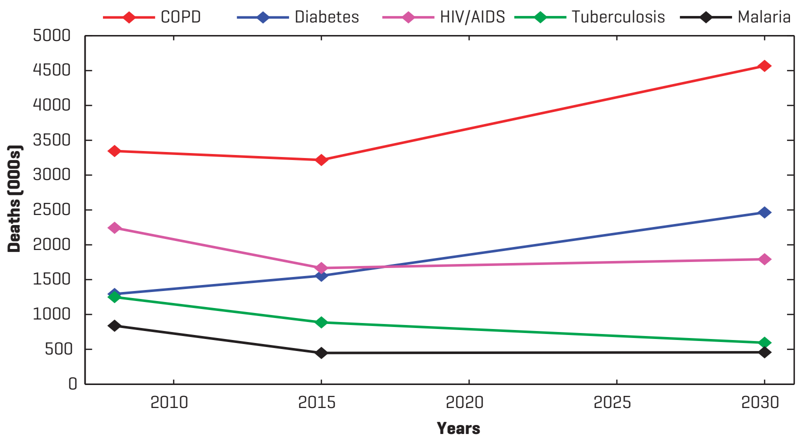

chronic patients often fear exercise if not under therapist supervision given the potential consequences of incorrect exercise techniques. However distance and cost often inhibit patients from attending a rehabilitation center regularly, especially in developing countries where COPD prevalence is higher.

While therapists can recommend daily exercises as "homework," both therapists and patients have currently no means to assess exercise performance during independent training. It is therefore essential to develop new systems and service concepts that permit chronic disease management at home. The use of smartphones, ubiquitous sensors, and network technology in healthcare systems could enable patients to perform additional physical training on their own, in addition to supervised training with a therapist.

During rehabilitation exercise, different errors can occur at the same time and should be identified accordingly. It is also essential to provide an error estimation algorithm that can handle different exercises with minimal adjustments to support training variety. Analyzing exercise performance is usually done by means of cameras—depth cameras and optical motion capture systems in combination with passive markers. In general, vision-based systems allow users to extract a human skeleton automatically, but require constrained environments to install and calibrate cameras. Due to these limitations, error-monitoring approaches started focusing on individual exercises or specific wearable training devices that helped to stratify error conditions. Various ambient and on-body device developments identified opportunities for continuous training and coaching in fitness and sports outside the lab setting. Often these approaches relied on multi-sensor information and pattern recognition methods, requiring individual learning of motion-pattern models. Although wearing multiple on-body sensors could provide high feedback accuracies, their cost and handling is challenging for patients. Smartphones, on the contrary, provide several integrated sensors to analyze data in real time and provide train-

Figure 2. Mortality rates due to COPD in different parts of the world, numbers in '000 (data from Lopez, A.et.al. 2006).

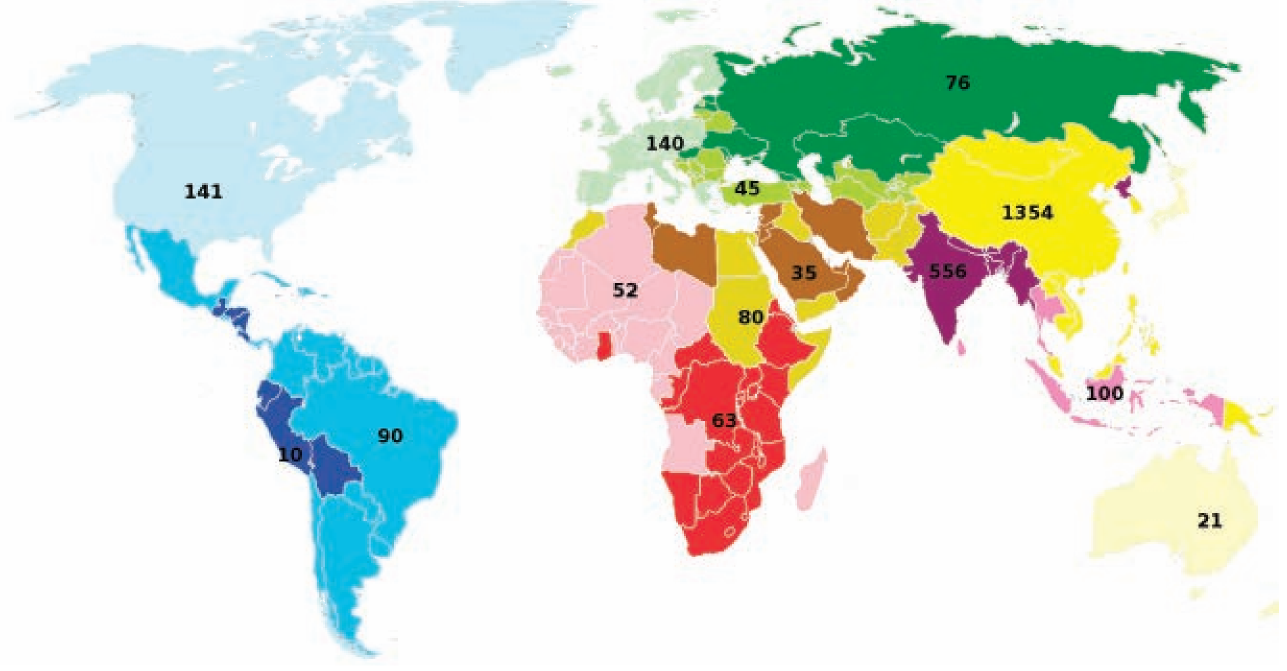

ing performance feedback. To minimize costs and other entry hurdles to personal rehabilitation training, smartphones may be the answer.

SMARTPHONE-BASED TRAINING APPROACH

COPDTrainer is a new smartphone-supported training application that considers the aspects mentioned previously and integrates into the usual clinical rehabilitation routine [3]. For COPDTrainer, a smartphone serves as a single measurement, estimation, and feedback device for assessing patient exercise performances. Recognition performance was evaluated for classifying execution errors, which is necessary to deploy the system in practice and especially in a clinical application. In this setting, the ability to perform particular motion exercises differs between trainees, due to individual motion constraints. Chronic patients, who often suffer from pathologies and muscle weakness, may not be able to perform exercises at the same speed or range of motion as another trainee. To overcome this problem the training approach adopted by COPDTrainer includes Teach and Train-modes as illustrated in Figure 3.

The Teach-mode allows therapists to personalize the system for a trainee under direct supervision. For example, during the regular physiotherapy practicing times any selectable exercise can be performed and the trainee learns from the therapist how to attach the phone and perform a particular exercise. Once an exercise is selected, illustrations are shown on the screen to remind the patient about the exercise execution. In Teach-mode, the therapist initially guides the patient during the first trials to perform the exercise accurately. Teach-mode recording begins once a large button on the phone's screen is pressed. A preset number of exercise repetitions (10 by default) will then be acquired from the phone's inertial sensors. From the recorded data, all necessary exercise model parameters, such as mean and variance of the duration and the range of motion of the limb during the 10 repetitions, are estimated and stored for further use during Train-mode. The derived parameters are shown on the smartphone, so the therapist and trainee can review them. If the therapist concludes the trainee did not perform the exercise with sufficient quality, the session could be repeated. Moreover, the system checks consistency of the exercise repetitions and can reject a Teach-mode session that shows extensive execution variability. These choices consider the regular clinical routines, where therapists have only 30 to 45 minutes per patient for assessment, therapy, and exercise training. Thus, complex interactions with the device were avoided.

It is therefore essential to develop new systems and service concepts that permit chronic disease management at home.

During Train-mode, the derived exercise models are arranged in a "to-do list" for the trainee to complete. This mode is intended for use by the trainee to exercise without therapist supervision, at the rehab center or at home. After selecting an exercise to be performed and starting the Train-mode, inertial motion data is recorded from the phone's sensors and processed in real time to count the exercise repetitions and detect errors. While training, COPDTrainer will provide acoustic feedback on the counted repetitions and notify when errors occur. For example, if the trainee practiced an exercise with the therapist before but starts to perform repetitions faster than during the Teach-mode, the system will provide the feedback "move slower." This feedback could prevent injuries from repetitive erroneous movements. Finally, after the configured number of repetitions is detected, the system will ask the trainee to stop and displays a summary of the execution performance.

Based on the observation that many fitness exercises have a repetitive structure, from training with free weights to cardio fitness motion, a sinusoidal motion model was considered. This method was chosen over others for two reasons: (1) Using machine-learning techniques requires a training set to obtain the classifier model. In particular, a sufficient number of exercise error instances would be required, but it is not feasible to let patients perform exercise errors due to the risk of injuries. (2) With machine-learning techniques, it is difficult to differentiate variations in performance of the same exercise from

Figure 3. COPDTrainer training approach.

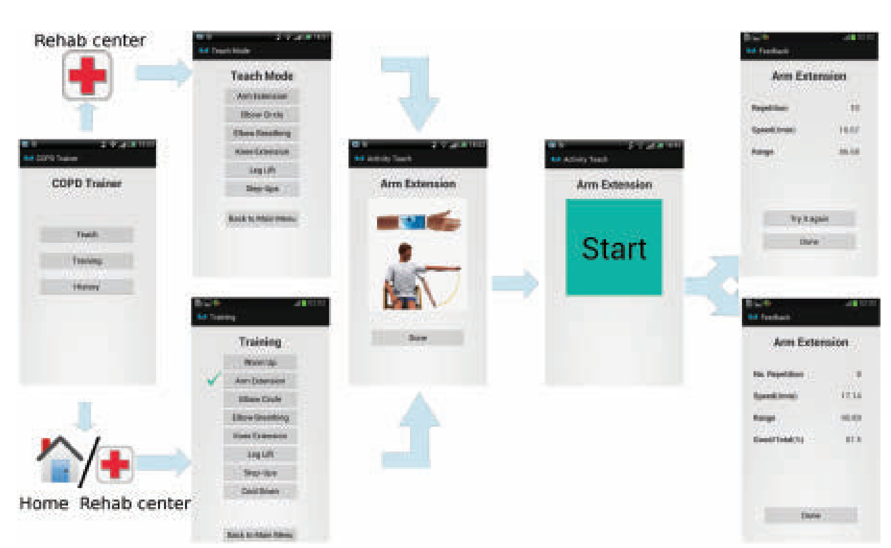

execution errors. Hence, error classes were formalized by considering deviations from the correct execution using the sinusoidal model.

For each exercise, a therapist or expert could choose a representative motion feature that represents a sinusoidal pattern. The feature can be based on a single raw axis of acceleration, gyroscope, and magnetic field sensor, or fused from several sensors of the phone, such as orientation estimates. For example, in a lateral arm abduction exercise, where the phone is attached to the wrist, the anterior-posterior orientation angle could be used as motion feature. The smartphone position at the body and feature need to be selected only once per exercise type. Exercises could be shared between patients, therapists, and clinics subsequently. Since the Teach-mode is performed under therapist supervision, no real-time feedback will be provided. Once the trainee completes an exercise session with a preset number of repetitions, the application loads the stored data and extracts the exercise model parameters.

The selected motion feature was filtered using a moving average to remove tremor-induced noise and sensor noise. The window size was set proportional to the amount of data acquired. This approach provided consistent results across different exercises. Since the number of repetitions is preconfigured, it was assumed that the total data amount recorded is proportional to the movement speed during the exercise execution: When a trainee performs the exercise faster, muscular tremor is lower, and thus, data averaging is reduced. Bounds were applied to the averaging window size to prevent ineffective averaging for very fast and slow repetitions.

By estimating the position of positive and negative peaks in the filtered motion feature, exercise repetitions were counted. For the arm abduction exercise, the selected feature is maximal when the arm is raised to shoulder height. It reaches its minimum value when the arm returns to the neutral position (arm aligned to the trunk). An adaptive, hill-climbing algorithm was then used to detect positive and negative peaks, given a starting peak threshold. While there are many alternatives, such as simulated annealing or tabu search, hill-climbing can achieve sufficient or better results if runtime is constrained, such as in the real-time system targeted here.

The use of smartphones, ubiquitous sensors, and network technology in healthcare systems could enable patients to perform additional physical training on their own.

The detection of local maxima and minima remains susceptible to detecting additional peaks (insertion), e.g. during vibrations, or to missing peaks (deletion) if the signal amplitude decreases. In situations where there was one insertion or deletion error in sequence, the alternating order of positive and negative peaks was interrupted. If two consecutive positive or negative peaks were derived, a peak correction algorithm was applied. This peak correction works by first removing redundant peaks and then inserting missing peaks that were missed during the first iteration of the hill-climbing algorithm. Time intervals between two consecutive peaks were also used to determine if there could be peaks missing. After segmenting the signal in single repetitions the following five parameters were derived: number of repetitions, mean and standard deviation of repetition duration, and mean and standard deviation of the range of motion. The repetition duration was derived from the time interval between two adjacent minima. The range of motion was derived from the magnitude difference between adjacent negative and positive peaks. The number of repetitions could be obtained by counting the number of maxima.

In contrast to the Teach-mode, Train-Mode operation requires online period estimation and subsequent performance analysis. A sliding window was used to segment the incoming data stream. The sliding window size was set to cover two average repetitions based on the parameters estimated in the Teach-mode, with an overlap of 75 percent between consecutive windows. The overlap ensures timely feedback during a newly detected repetition. To provide timely feedback, i.e. before the trainee starts a subsequent repetition, the first half of a repetition was evaluated to estimate duration and range of motion estimates. In preliminary tests, we observed the error incurred by considering only half of a repetition was negligible. The derived duration and range of motion estimates were used to compare with the parameters estimated in the Teach-mode. For duration and range of motion, each repetition performance is estimated based on a Gaussian distribution. The performance of each exercise is classified into three class types: in-between, under, and above the ranges. In total, we considered nine different classes to which each performed exercise repetition could be associated. After the performance class corresponding to the ongoing repetition was evaluated, an audio feedback was provided to the trainee, who was notified if a repetition was erroneously performed.

COPDTRAINER EVALUATION

Advised by three therapists, and after consulting COPD guidelines, speed of motion (corresponding to the period frequency) and range of motion (corresponding to the feature amplitude) were derived from the sinusoidal pattern of each exercise repetition. In kinesiology speed and range of motion, together with their relative tolerances and the number of repetitions, are considered standard measures for exercise monitoring. Estimating movement speed during exercises is useful to educate patients in breathing techniques (i.e. by exercising the patient can learn how to breathe with correct timing). Based on these exer-

Figure 4. Performance classes, feedback, and condition used to identify exercise quality.

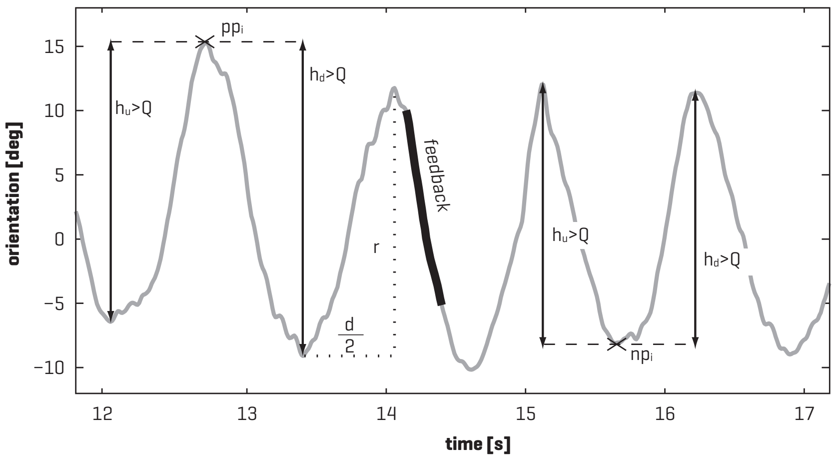

Figure 5. Illustration of the exercises selected for the training system evaluation. The patient is wearing the smartphone (red circle) on limbs that are involved in the different exercises.

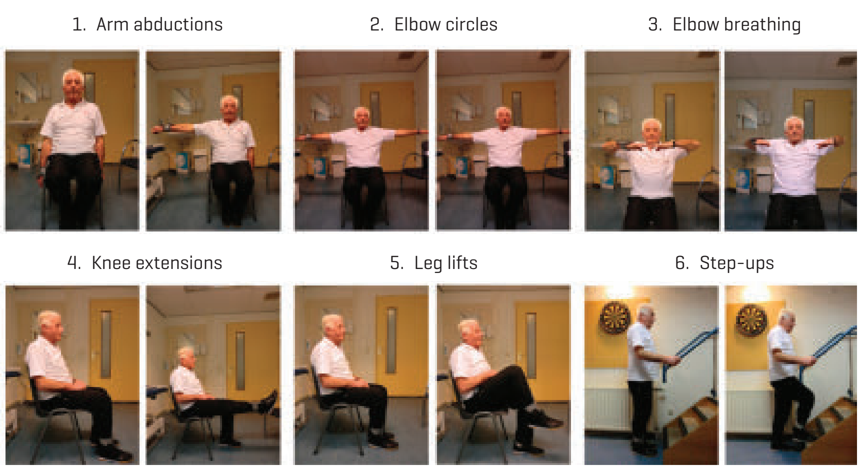

cise quality parameters, it is possible to derive performance classes, such that the classes are applicable to various exercises performed by repetitive movements. During the Teach-mode, exercise repetitions are used to represent repetition range and duration parameters using two normal distributions. In the Train-mode, these model parameters are used to identify nine performance classes and can be seen in Figure 4. Six exercises, shown in Figure 5, were chosen for daily training at home according to the COPD guidelines and in consultation with therapists. The exercise set consisted of three upper limb muscle exercises: arm abductions (AA), elbow circles (EC), and elbow breathing (EB); and three lower body muscle exercises: knee extensions (KE), leg lifts (LL), and step-ups (SU).

To test and evaluate the training system, two sets of experiments were conducted. Initially, the system was validated with healthy participants using a scripted protocol, where all performance classes have been equally represented. Subsequently, the training system was evaluated in an intervention study with COPD patients performing normal therapy training sessions. The validation with healthy participants showed an overall accuracy of 96.2 percent. The intervention study with seven COPD patients showed a trainee performance classification rate of 87.5 percent, while repetitions were counted at 96.7 percent accuracy.

Based on those results, we concluded a smartphone-based training system can be used to assess the performance and execution quality of a rehabilitation exercises in COPD patients. Based on the system performance and feedback efficacy, we believe our approach and developed methods will be a vital basis for future investigations on training systems for different patient groups. Additional steps are needed to confirm the clinical relevance and integration into clinical practice. In this regard, we consider this work as a pilot study, providing the basis for validating COPDTrainer in a clinically supervised intervention at the patient's home. We hope the COPDTrainer application will become an everyday tool for patients to improve and maintain their health state.

References

[1] Lopez, A., Mathers,C., Ezzati, M., Jamison,D., and Murray, C. Global and regional burden of disease and risk factors, 2001: Systematic analysis of population health data. *The Lancet* 367, 9524 (2006),1747-1757.

[2] Buist, A. et. al. International variation in the prevalence of COPD (the bold study): A population-based prevalence study. *The Lancet* 370, 9589 (2007), 741-750.

[3] Spina, G., Huang, G., Vaes, A., Spruit, M., and Amft, O. COPDTrainer: A smartphone-based motion rehabilitation training system with real-time acoustic feedback. In *Proceedings of the 2013 ACM Internetional Joint Conference on Pervasive and Ubiquitous Computing* (Zurich, Sept. 8-12). ACM Press, New York, 2013.

[4] Lopez, A.et.al. Chronic obstructive pulmonary disease: Current burden and future projections. *The European Respiratory Journal 27*, 2 (2006).

Biographies

Gabriele Spina received B.S and M.S. degrees in biomedical engineering from Università Campus Bio-Medico in Rome. Currently he is a Ph.D. candidate in the ACTLab research group at TU Eindhoven. His main research focuses in the use of emerging technologies (mobile, ubiquitous sensor and network technology) in healthcare systems to monitor patient's status and provide insight into daily life activities.

Oliver Amft is an assistant professor at TU Eindhoven, where he heads the ACTLab research group, and a senior research advisor at the Wearable Computing Lab, ETH Zurich. His research focuses on multi-modal activity recognition and human behavior inference algorithms for ubiquitous sensing systems, with applications (among others) in healthcare, sports, and building automation.

Capturing Human Motion One Step at a Time

The design, construction, and deployment of a pressure-enhanced IMU system that fits in the bottom of your shoe.

By Rolf Adelsberger

DOI: 10.1145/2538692

Digitizing human motion can provide deeper insight into both the physical and mental properties of a subject. Once captured, the motion of different body parts can be quantified, analyzed, and related to one another to extract useful information. In medicine, doctors try to augment diagnoses, or even base new diagnoses, on objective motion data. Sports and biomechanics focus on objective measures of physical properties from motion data, e.g., angles, accelerations, rotations, and torque. In this article, we provide a focused view on the development and deployment of our motion-sensing device and discuss our use of the system in both a medical setting and in experiments with athletes.

INTRODUCTION

We aimed to create a wearable sensor system that tells us more about mobility than a single accelerometer or pedometer, but is as unobtrusive as possible. Initially, we wanted to assess data to determine gait patterns, fitness, or mobility of elderly people. The fitness of a subject is defined depending on context: The time required to run 100 meters might be used as a fitness indicator for an athlete, whereas a mobility index (MI) is often employed for elderly people.

There are an abundant number of mobility indices used in practice. The Barthel-Index, for example, captures everyday activities like eating and drinking, but also more specific activities like climbing stairs and sitting on chairs. A more focused and widely used index is the timed-up-and-go (TUG) test: The time required for a subject to rise from a seated position in a chair and to walk five meters is measured and compared with a normative score [1]. The longer it takes the person, the lower the score. Prior studies have shown TUG correlates very well with a subject's risk of falling. Since injuries caused by falls are usually more severe for elderly people, falls should be prevented as often as possible.

Motivated by prior research [2], we knew important information could be obtained from analysis of temporal gait patterns in (elderly) subjects: Step frequency, stance time, swing time, anterior-posterior, and medio-lateral sway, are important statistical features affected by a subject's fitness or mobility. These features can be captured with inertial measurement units (IMUs) comprised of an accelerometer, a gyroscope, and often a magnetometer. On the other hand, balance, posture, TUG scores, center-of-pressure characteristics, are also established measures that assess different parts of a subject's mobility. These features could be estimated using an IMU-based system, but not possible if you only use a single-sensor device. After evaluating possible existing alternatives (multi-IMU systems, optical systems, pressure sensitive flooring or device, etc.) we decided to create our own device: A pressure-enhanced IMU system that can be worn unobtrusively in a shoe.

A NOVEL DEVICE

Our first task was identifying key features. Unobtrusiveness was an important design goal. This requires small packaging and no wires; it should be unnoticeable to the user. Our lab had previously developed a small IMU sensor with a wireless communication channel; the missing piece was the pressure-sensitive part. There are pressure-sampling systems that use thin polymer foils and incorporate

Image by Wave Break Media


92
117
122
87
-279cal
132
79
135
147
87

Figure 1. The PIMU device showing the electronics with casing and the pressure insole.

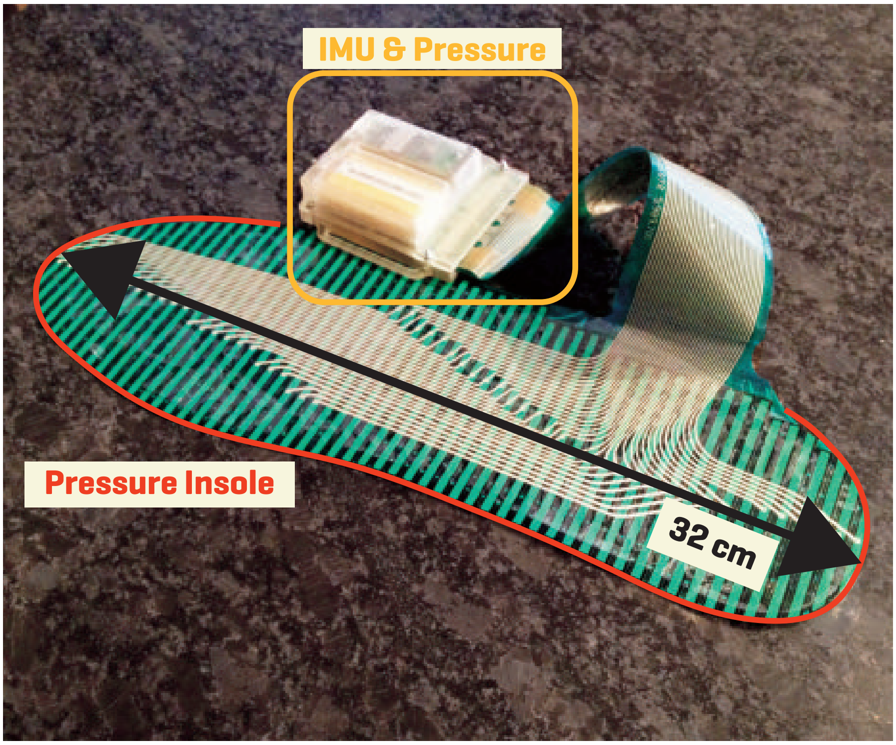

pressure-sensitive electronics. For example, TekScan builds tethered systems where the sampling circuitry communicates to an aggregator device attached to a subject's waist. Via a wired connection to a base station, the pressure insoles can be sampled at about 100 Hz.

We concluded that a combination of both devices—an IMU and a pressure-sensitive insole—would fit our needs perfectly. Not only would such a device assess temporal features from steps of a subject, but the pressure insole would allow assessment of a high-resolution time series of pressure maps. Tests, like TUG, could be performed automatically by detecting situations of standing, sitting, and walking, and of estimating distance walked. With wireless communication, not only are per-foot center-of-pressure (COP) calculations possible, but a body-COP assessment is also feasible. We decided to call our device PIMU, for pressure-sensing IMU.

The layout and population of the circuit board turned out to be straight-forward. The pressure-sensitive insole was constructed in a matrix-style layout: On one side the force-sensing resistors (FSR), which are the sensor elements, are connected in a column order; on the other side of the foil, there is a row-like connection. The FSRs change their resistance relative to the applied pressure. In a nutshell, we connected the pins of the central processing, which could be configured as outputs (e.g., the columns) to one side of the sensor foil. The other side was connected to the rows of the sensor foil. At runtime, the sensor board enables one output pin and then samples the voltage for all input pins. A complete set of pressure readings, i.e., one sole, is sampled if every output pin was once assigned high.

The remaining obstacle was enabling communication between the IMU and the pressure system. There are multiple communication standards for electronics; a very prominent one is the inter-integrated circuit, or I^2C. The processor on the IMU already uses I^2C to talk to the sensors. Since it is a bus system, adding a new party to the communication was straightforward. The only things needed were a wired connection from the pressure system to the I^2C wires of the IMU and software adaptions on the inertial sensor in order to enable communication with the other system.

Without going into detail, both operating systems on the two sensors are designed in an interrupt-driven philosophy. That means instead of polling for answers from external sensors or communication chips, the processors are notified whenever something new happens). This design principle reduced power consumption on both systems by more than 60 percent and relaxed the system load so complex calculations could be implemented in real time on the sensor boards. Our system can run continuously for more than 12 hours sampling IMU data at 128 Hz and pressure data at 100 Hz.

TURNING MULTIPLE MODULES INTO ONE SENSOR SYSTEM

We used a 3D-printer to print housings receiving both sensor modules, a battery, and the insole-connector in a small volume. PIMU incorporates an ANT+ enabled wireless communication chip. It is compatible with heart-rate belts and similar sports equipment by Garmin and other manufacturers. ANT is a very low-power communication protocol and is implemented by specific types of smart phones. Additionally, an ANT USB dongle enables this communication channel on a regular PC. We decided to use a smartphone to control the sensors and display sensor data. The application runs on Android OS and enables the operator to start and stop the sensors, display real-time data, and configure sampling parameters.

Often in a multi-device setup, good inter-device synchronization is difficult or even unattainable and has to be performed offline. Our setup allows us to have a very accurate synchronization between any two sensors; the difference between two sensors at the beginning of sampling is at most four microseconds. We achieved this high accuracy with a carefully designed interrupt mechanism in the sensor's operating system. The communication interrupts are assigned a high priority within an interrupt context. Only if a sensor is currently in another interrupt context could there be a delay (at most four microseconds). The clocks on the sensor boards exhibit a maximal drift of 10 parts per million. In our configuration, this results in a maximal drift between any two sensors of about 36 milliseconds over a 60-minute time period.

DEPLOYMENT

Once the hardware was built and wired and the operating system and back-end applications were programmed, the system was ready to be used. Presented here are increasingly complex tasks that were performed by our system in deployment.

Gait analysis. In geriatric medicine, it is well known that gait performance decreases with increasing cognitive load for many elderly people. Hence, affected subjects tend to walk more irregularly and have more difficulties maintaining a straight path when asked to perform a cognitive task. However, it has been hypothesized that mental training, as well as physical training, could reduce the impact of cognitive load on gait patterns. Consequently, this would also lower the risk of falling. For a study evaluating training effects on elderly people, we used both PIMUs and commercial IMUs that were attached to the subjects' legs while they were walking on a treadmill. For each of the 16 subjects, we recorded three sessions: The first prior to any training, the second to assess performance in the middle of training, and the last once training was complete. The training period occurred over 14 weeks. Our results showed with high statistical significance that training is beneficial to the gait performance of elderly people. Furthermore, even training just once a week improves overall gait performance and reduces possible impacts of increased cognitive load on gait performance [4]. Our gait analysis showed the IMU part of our system (i.e., no pressure data) is useful for the assessment of slow movements.

Athlete-centric motion analysis. We next decided to tackle a more challenging problem: Exploiting the high-accuracy synchronization between any two PIMU sensors. For this purpose, we took a scientific detour into a new domain: weight lifting. In many sports, the synchronization between different body parts is of high importance for good performance. This is especially true for weight lifting, where a very heavy mass, e.g., a barbell, is moved externally to the body. Weight-lifting athletes practice exploiting the momentum they created in every phase of a movement. For example, a barbell is initially accelerated upwards with lower-body muscles until it reaches a certain height, at which point the athlete tries to switch as smoothly as possible to upper-body muscles.

We used our system to test whether beginner athletes are easily distinguishable from more experienced athletes by looking at the synchronization between lower-body and upper-body movements. We asked 12 athletes of different experience levels to wear our sensor system while they performed front squats with a consecutive overhead press (an exercise called a "thruster"). We were very pleased with the results. Analyzing the power generated by hip and arm movements measured using our PIMU device, enabled us to to classify individual athletes by experience level with more than 90 percent accuracy.

Balance assessment. So far, we validated our PIMU component for performing motion analysis in two different settings. Our main focus, however, is posture and stability anal-

Figure 2: An athlete wearing three PIMUs in the start position of a "thruster."

ysis of elderly patients. In more recent studies, we used our system to assess postural stability of patients with balance deficiencies.

There are uncountable medical causes that eventually lead to difficulties maintaining a balanced posture. The human balance system comprises three main components: proprioception (the sense of the relative position of neighboring parts of the body), the visual system, and the vestibular system. If any of these body systems are affected by a medical condition, the overall sense of balance (equilibrioception) can be reduced.

In collaboration with a local hospital, we attended several functional gait analysis (FGA) sessions. This analysis screens patients to determine if they need physiotherapy to improve their posture stability. The sessions seldom incorporate technical tools; usually, a medical doctor or physiotherapis tasks patients to perform several items from a test battery and estimates the score on an ordinal scale. This assessment is frequently difficult even for well-experienced experts with 10-plus years of practice. We wanted to create a tool for medical experts to better—objectively and quantitatively—assess the patient's movements.

Typical tests in an FGA include a) walking on a line with closed eyes and b) standing heels together with eyes closed. The former is rated according to an estimated maximal drift from the optimal line. The latter would be harder to assess, because an objective statement about a patient's stability can only be made if a patient needs to take a step. However, stability is not a binary entity; it can be modeled as a function of the coordinates of a subject's center of mass (COM). Measuring the COM and its projection on the ground can give an accurate assessment of the subject's balance. If that point falls inside the base support spanned by a person's feet, his or her posture is stable. Assessing the COM requires accurate knowledge of the mass distribution of a subject's limbs, as well as the ability to track the trunk and all limbs with high accuracy. Optical motion capture systems are very good at that task. Measuring the COP

Figure 3: Pressure data from PIMU device.

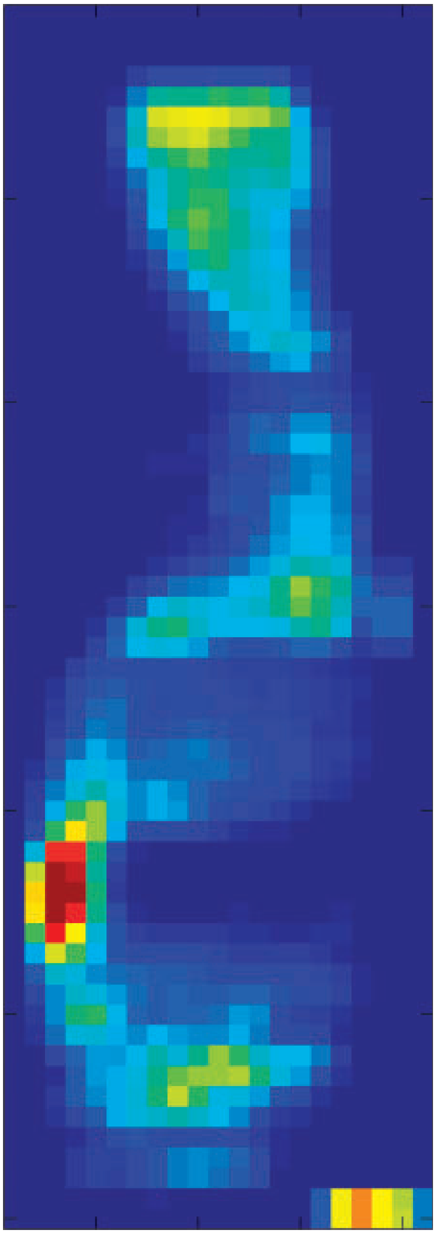

does not unveil the distribution of mass, but it does tell us about the applied forces. These forces result from body movement initiated by a subject trying to maintain balanced posture.

Prior research has shown by analyzing the time series of COP displacements for a standing subject, the stability of the subject can be estimated very accurately without knowing the location of the COM. An intuitive explanation could be a stable person shows less sway in his or her COP than a subject fighting for stability. The technical term for that proxy is stabilogram diffusion analysis, or SDA [6].

We implemented algorithms on the PIMU sensors that track the COP of a subject in real-time. For validation, the performance of our device was compared to that of a medical treadmill with incorporated pressure sensors. The comparison proved our system is a valid alternative to a static COP-assessment system. For those FGA items assessing static posture stability, our system can report an objective measure to a medical expert, quantitatively aiding in score assignment.

CONCLUSIONS AND OUTLOOK

Human motion is an amazing source of information for answering many questions in sports and medicine, as well as in everyday life. In our journey designing the PIMU system—from the first step of problem identification, to the device's design, creation, and, finally, deployment—our results have shown that even a single device can enable multifaceted, motion analysis research in a variety of different problem domains.

References

[1] Podsiadlo D. and Richardson, S. The timed "up and go": A test of basic functional mobility for frail elderly persons. *Journal of the American Geriatrics Society 39*, 2 (1991).

[2] Hausdorff, J., Schweiger, A., Herman, T., Yogev-Seligmann, G., and Giladi, N. Dual-task decrements in gait: Contributing factors among healthy older adults. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences* 63, 12 (2008).

[3] Adelsberger, R. and Tröster, G. PIMU: A wireless pressure-sensing IMU. In *IEEE Proceedings of the 8th International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP)*, Apr. 2013.

[4] Adelsberger, R., Theill, N., Schumacher, V., Arnrich, B. and Tröster, G. One IMU is suffcient: A study evaluating effects of dual-tasks on gait in elderly people. In *MobiHealth*. Springer, New York, 2012.

[5] Adelsberger, R. and Tröster, G. Experts Lift Differently: Classification of Weight-lifting Athletes. In *IEEE International Conference on Body Sensor Networks (BSN)* (Cambridge, MA, May 6-9). 2013.

[6] Collins, J. and Luca, C. Open-loop and closed-loop control of posture: A random-walk analysis of center-of-pressure trajectories. *Experimental Brain Research* 95, 308 (1993).

Biography

Rolf Adelsberger has a M.Sc. in computer science from ETH Zurich. His first steps in research were in the area of computer graphics where he was looking into motion capture, 3-D imaging / 3-D video and related projects. He is currently pursuing a Ph.D. in electrical engineering and works with the wearable computer lab at ETH.

mHealth @ UAH: Computing infrastructure for mobile health and wellness monitoring

New health care systems that integrate wearable sensors, personal devices, and servers promise to fundamentally change the way health care services are delivered and used.

By Mladen Milosevic, Aleksandar Milenkovic, and Emil Jovanov

DOI: 10.1145/2539269

Mobile health (mHealth) represents the use of mobile wireless communication devices to improve health outcomes, healthcare services, and health research [1]. mHealth monitoring systems typically integrate wearable physiological sensors, personal devices like smartphones, and servers accessed over the Internet. They have emerged as a promising technology for real-time, unobtrusive, and continuous health and wellness monitoring of individuals during activities of daily living. Such systems promise to radically modernize and change the way healthcare services are deployed and delivered. They allow an individual to closely monitor changes in his or her vital signs and provide feedback to help maintain an optimal health and wellness status. When integrated with healthcare providers, these systems can even alert medical personnel when life-threatening changes occur. In addition, mHealth monitoring systems can be used for health monitoring of patients in ambulatory settings: as part of a diagnostic procedure, an optimal maintenance of a chronic condition, or a supervised recovery from an acute event or surgical procedure. They can also be used to monitor adherence to treatment guidelines (e.g., regular cardiovascular exercise) or to monitor effects of drug therapy.

At the University of Alabama in Huntsville (UAH) an mHealth infrastructure, including both hardware and software components, was created to support research and education in the area of computer systems for mobile health and wellness monitoring. It is designed to help address critical design issues in the next generation of health monitoring systems—including their functionality, reliability, and energy-efficiency—to support creation of a repository with vital signs and physical activity parameters during normal daily activities, and to enable rapid prototyping of new monitoring applications.

HEALTH AT THE TOUCH OF A BUTTON

Convergence of smart biosensors, smartphones, and cloud computing services have enabled the development and proliferation of affordable mHealth monitoring systems capable of continuous health and wellness monitoring. Advances in sensor

technology have enabled miniature smart sensors to unobtrusively monitor physiological signals, body posture, type and level of physical activity, and environmental conditions. Physiological signals include heart electrical activity (electrocardiogram /ECG), muscle electrical activity (electromyography/EMG), brain electrical activity (electroencephalography/EEG), pulse and blood oxygen saturation (photoplethysmography/PPG), blood pressure, respiration/breathing rate, galvanic skin response (GSR), blood glucose level, and body temperature. In addition to the physiological signals, mHealth wearable monitors may include sensors that can help determine the user's location, discriminate between user's states (e.g., laying, sitting, walking, running), or sensors that can help estimate the type and level of the user's physical activity (e.g., low-, moderate-, or high-intensity aerobic activity). Since environmental conditions may influence the user's physiological state or accuracy of the sensors, mHealth monitors may integrate information about environmental conditions, such as: humidity, light, ambient temperature, atmospheric pressure, and noise.

Figure 1. Data flow in mHealth's three-tiered architecture.

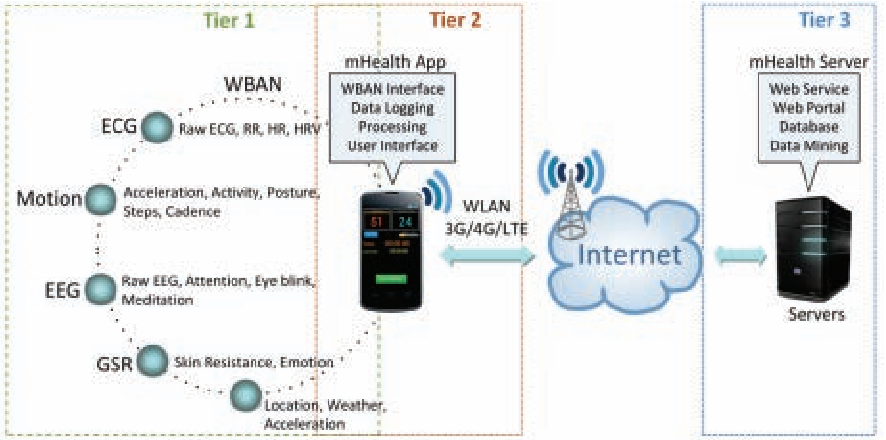

Availability, affordability, and excellent performance make smartphones an ideal platform for mHealth applications. According to a report from August 2013, 225 million smartphones were sold worldwide in the second quarter of 2013, which represent an increase of 46.5 percent compared to the same period in 2012 [2]. With the recent proliferation of smartphones and tablet computers, the number of health monitoring and wellness applications has exponentially increased. According to a report from March 2013, more than 97,000 mHealth applications are listed on a variety of application stores [3]. Moreover, Google and Apple recognized this trend and made modifications in their operating systems to directly support health and wellness applications. The Android operating system incorporates a

Photograph by Warren Goldswain

service that detects the user's current physical activity, such as walking, driving, or standing still. Apple went one step further with the latest iPhone 5S by designing and implementing a separate motion coprocessor to analyze user's activity from the motion sensors (accelerometer, gyroscope, and magnetometer). The availability of affordable smartphones and wearable devices, their widespread use, and consumer acceptance create new opportunities for users and healthcare professionals. An increasing number of users, who actively monitor their own health and fitness status, further underscores this trend [4].

MHEALTH @ UAH

The mHealth infrastructure at UAH is designed as a three-tiered architecture with wireless body area sensor networks and other physiological monitors at Tier 1, personal computing devices at Tier 2, and mHealth servers at Tier 3. This is represented in Figure 1.

Tier 1 consists of one or more body area networks (BANs) or body sensor networks (BSN) optimized for a specific health monitoring application. Each network integrates one or more wearable and intelligent sensor nodes. We rely on commercially available sensors and wearable monitors that sense vital signs, body posture, and environmental conditions. They range from inexpensive sensors (less than $100) intended for fitness monitoring applications to more sophisticated monitors designed for research (more than $2,000).

For monitoring cardiac activity we use a range of monitors differing in

Figure 2. mHealth @ UAH.

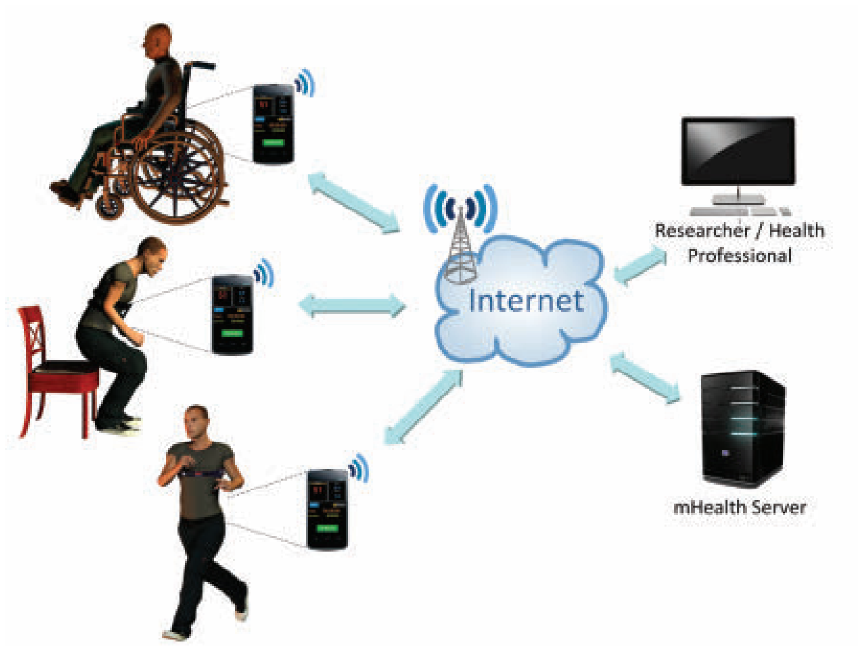

Figure 3. iTUG test phases and smartphone instrumentation of the subject.

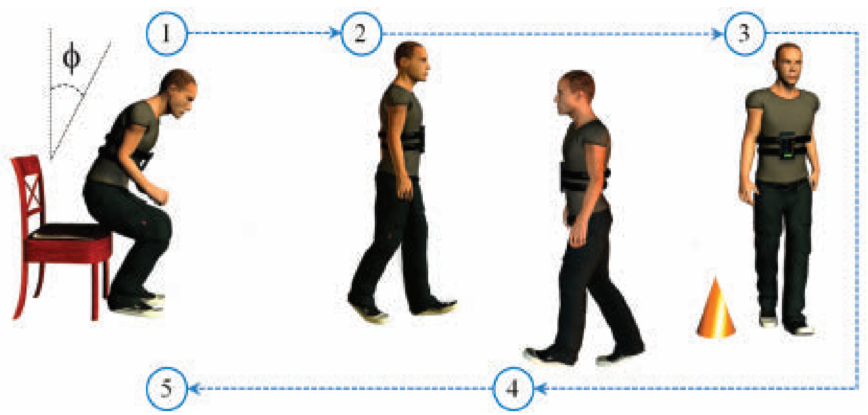

Figure 4. sTUG: Smartphone TUG Android application screen displaying the parameters of the TUG test.

form factor, weight, functionality, accuracy, and cost. They range from fitness grade monitors that can report only an average heart rate to medical grade monitors that can report and record interbeat intervals (RR intervals) and electrocardiogram (ECG). For example, the Garmin ANT+ or Zephyr HxM heart rate monitors are a good choice for applications that have a long battery life and a small form factor as prime requirements. The Zephyr BioHarness 3 and Hidalgo Equivital 2 physiological monitors, capable of recording RR intervals and raw ECG signals, are a good choice for applications where accuracy and resolution are prime requirements. They also include additional sensors such as a three-axis accelerometer and a respiration sensor, and use the Bluetooth wireless interface for communication at Tier 2.

For monitoring brain electrical activity we use Zeo sleep monitors, NeoroSky MindSet EEG sensors, and Emotiv EEG neuoroheadsets. The Zeo sleep monitor is a low-power headband with a single channel EEG intended for sleep studies. The MindSet EEG provides a single channel EEG in the form of a wireless headset, whereas the Emotiv EEG headset offers 14 channels of EEG sampled, filtered, and reported through Bluetooth to a custom application.

For monitoring physical activity, body posture, and transitions we use a range of commercially available sensors, such as the Garmin ANT+ foot pod sensor and the Garmin ANT+ bike sensor or inertial sensors featuring accelerometers, gyroscopes, and magnetic sensors. The foot pod sensor measures the number of steps made and speed during walking/running, while the speed/cadence sensor measures cycling speed. Both sensors use the low-power ANT+ wireless interface for communication at Tier 2.

Personal devices like smartphones can also be utilized for sensing body posture, physical activity, and environmental conditions. For example, a Google Nexus 4 smartphone includes a three-axis accelerometer, a three-axis gyroscope, a three-axis magnetometer, a barometer, a proximity sensor, an ambient light sensor, a GPS (Global Positioning System), and two cameras.

Personal applications running on a personal device (e.g., Android and Apple/iOS smartphones, tablets, or personal computers) represent Tier 2 of the proposed architecture. Applications are designed to facilitate (1) interface and management of a variety of sensors in the sensor network; (2) data retrieval from individual sensors, data logging, and analysis to extract health status information; and (3) user interface providing real-time feedback with health parameters and recommendations (e.g., guided rehabilitation or exercise). The collected health status information is periodically uploaded to the mHealth servers over the Internet. The majority of applications are developed for Google's Android and Microsoft Windows operating systems.

A group of servers providing storage, access, visualization, and support for data mining of physiological records forms Tier 3 of the mHealth infrastructure. The servers are running a free operating system, Ubuntu Server, and are designed and implemented to work as virtual machine appliances in either open source VM VirtualBox or proprietary VMWare environment. This approach offers flexibility and easy deployment and migration to new physical platforms or even to cloud infrastructure. Tier 3 of the mHealth infrastructure is composed of three main components: mHealth Database, Web API, and Web Portal. System architecture and sample applications are presented in Figure 2.

The mHealth database is developed using Oracle's MySQL relational database. The open-source database is specifically designed to support efficient storage of a variety of physiological records and record annotations. Each record has information about the subject, equipment used to collect records, and conditions under which the data are recorded. Physiological records can be organized by application type, and each record is precisely time-stamped. In addition, the database provides support for management and guidance of a variety experiments in research environment. Experiments can be conducted using a specific protocol and have authorized investigators, a list of sessions with individual participants, and individual physiological, activity, and multimedia records.

The Web API component is designed to be an intermediary between the personal devices and the mHealth database. It accepts data from personal devices and stores it into the database, and also allows personal devices to retrieve stored data from the database. Any action using Web API requires successful authentication. Upon a successful authentication, a Web session is created, allowing further execution of Web API requests without additional authentication. After a predefined period of inactivity, the session automatically expires and the authentication process has to be repeated.

The Web Portal component provides easy access to physiological data and its basic visualization. It requires only a Web browser to access a recorded session in the mHealth database. It is developed using the Sencha JS framework. Each authenticated user is allowed to access only a subset of data he/she is authorized to access. The user can easily visualize data by selecting the desired session and the particular signals inside the session.

> **The availability of affordable smartphones and wearable devices and their widespread use and consumer acceptance create new opportunities for users and healthcare professionals.**

EXAMPLE APPLICATIONS

At UAH we originally developed two mHealth applications: sTUG and mWheelness. sTUG quantifies and automates a standard Timed-Up-and-Go (TUG) test used to assess mobility of individuals. mWheelness monitors physical activity of individuals who rely on wheelchairs for mobility.

Real-time quantification of TUG test. TUG is a frequently used clinical test for assessing balance, mobility, and fall risk in the elderly population and for people with Parkinson's disease. It is simple and easy to administer in an office, and thus can be used in screening protocols. The test measures the time a person takes to perform the following tasks: rise from a chair, walk three meters, turn around, walk back to the chair, and sit down. Longer TUG times have been associated with mobility impairments and increased fall risks. TUG duration is also sensitive to therapeutic interventions, such as in Parkinson's patients.

We have developed a smartphone application called sTUG that completely automates the instrumented Timed-Up-and-Go (iTUG) test so that it can be performed at home [5]. sTUG captures the subject's movements utilizing a smartphone's built-in accelerometer and gyroscope sensors, determines the beginning and the end of the test and quantifies its individual phases, and optionally uploads test descriptors into the mHealth server.

A subject mounts the smartphone on his/her chest or belt and starts the application, as illustrated in Figure 3. The application records and processes the signals from the smartphone's gyroscope and accelerometer sensors to extract the following parameters that quantify individual phases of the iTUG: (a) the total duration of the TUG, (b) the total duration of the sit-to-stand transition, and (c) the total duration of the stand-to-sit transition. In addition, we extract parameters that further quantify body movements during sit-to-stand and stand-to-sit transitions, including the duration of sub phases, maximum angular velocities, and upper trunk angles. These parameters are recorded on the smartphone and optionally uploaded to the mHealth server. The application stops monitoring auto-

Figure 5. Smartphone instrumentation of a wheelchair.

matically once it detects the end of the stand-to-sit transition. Figure 4 shows a report generated by the application at the end of a TUG test.

sTUG is developed for the Android operating systems and requires a smartphone with the accelerometer and gyroscope sensors running Android 2.3 or above. The application has been tested on a Nexus 4 smartphone, a Motorola RAZR M, and a RAZR HD.

We believe this application could be of great interest for older individuals and Parkinson's disease patients as well as for healthcare professionals. The procedure requires minimum setup (a chair and a marked distance of three meters) and inexpensive instrumentation (a smartphone running the sTUG application is placed on the chest or belt). The feedback is instantaneously provided to the user in a form of a report with the values of all significant parameters that characterize the TUG test. It is easy to use and users can take multiple tests in a single day at home (e.g., to assess the effects of drugs). With automatic updates to the mHealth server, caregivers and healthcare professionals can gain insights into overall wellness of the subjects. For example, they can assess the impact of therapeutic interventions (e.g., impact of drugs) by analyzing the parameters from multiple tests performed in a single day. Healthcare professionals and researchers can monitor and evaluate the evolution of a disease by analyzing the trends in the parameters collected over longer periods of time.

Real-time monitoring of activity of wheelchair users. Physically inactive individuals are almost twice as likely to develop coronary heart disease compared to those who exercise regularly. Recent estimates suggest the impact of physical inactivity on mortality risk is approaching that of tobacco as one of the leading causes of death in the able-bodied population. People with limited ambulatory skills who use wheelchairs for mobility are especially at high-risk for all inactivity-related diseases. For example, it has been reported that a person with a spinal cord injury (SCI) has a significantly greater risk of mortality from coronary heart disease (225 percent) than an able-bodied person. According to a 2005 U.S. Census Bureau's Survey, more than 3.3 million Americans use some type of wheelchair for mobility and with the aging population this number is likely to continue to grow.

In order to provide an affordable, reliable, and easy to use solution for monitoring the physical activity of users who rely on wheelchairs for mobility we developed a smart wheelchair [6]—a common wheelchair instrumented only with a smartphone that is used to track a user's physical activity. The system can record, log, display, and communicate information about the user's physical and heart activity during normal daily activities or exercise sessions. For monitoring the user's physical activity we utilized the smartphone's built-in sensors such as a magnetic sensor for monitoring wheelchair speed and distance traveled, an accelerometer for monitoring smartphone's orientation and wheelchair inclination, and a proximity sensor to determine whether the wheelchair is hand-propelled or pushed. In

Figure 6. mWheelness Android application screens.

addition, we employ a wearable chest belt to monitor and record the user's heart activity and energy expenditure. A smartphone application called mWheelness collects data from the sensors and performs periodic uploads to the mHealth server.

Figure 5 illustrates the proposed wheelchair instrumentation with a smartphone. The smartphone is placed in a holder on a side of the wheelchair. The smartphone's magnetic sensor senses the x, y, and z components of the magnetic field as illustrated in Figure 5. By placing a small magnet on the wheel, we induce a change in the magnetic field sensed by the magnetic sensor of the smartphone when the magnet moves over the smartphone. This change produces a characteristic signature in the magnetic field signals that can be sensed, recorded, and processed on the smartphone. By processing the magnetic field signals we can detect and timestamp an event, when the magnet moves right over the smartphone, which corresponds to one revolution of the wheelchair's wheel.

A smartphone's accelerometer measures proper acceleration and is typically used to keep the screen upright regardless of the smartphone orientation. In our setup we process the x, y, and z acceleration components to determine smartphone's orientation, i.e., whether it is placed in the wheelchair holder or not. Activity recording is enabled only when the smartphone is properly mounted on the wheelchair. In addition, the accelerometer data is used to determine slope of the wheelchair, which can further be used to determine vertical gain and loss during exercise.

A smartphone's proximity sensor is typically used to determine when the smartphone is brought up to the user's ear and usually acts as a binary sensor. In our deployment, the smartphone's proximity sensor is used to determine whether the user hand-propels the wheelchair or it is pushed. This information can be used to further qualify the user's activity.

Figure 6 shows one of the characteristic screens of mWheelness. The user starts recording physical activity and heart activity by pressing the start/stop recording button, although the processing of the signals from the magnetic sensor will not start before the smartphone is in the upright position. During an exercise session mWheelness displays current inclination, speed, and distance traveled. In addition, it displays information about heart activity.

Availability, affordability, and excellent performance make smartphones an ideal platform for mHealth applications.

The mWheelness application has been tested on several Android smartphones (Nexus 4, Motorola RAZR M, and HTC One X) in controlled and free-living conditions. The controlled experiments were conducted on a treadmill while varying speed and inclination. Distance traveled and inclination as reported by the application, were then compared against the corresponding parameters reported by the treadmill.

CONCLUSION

The infrastructure proved very effective in supporting research projects, course projects, and senior design projects in the exciting and emerging area of mobile health monitoring. More information about the mHealth infrastructure at UAH can be found at http://portal.mhealth.uah.edu. mHealth infrastructure developed and implemented at the University of Alabama in Huntsville was supported in part by NSF grant 1205439 *mHealth - Computing Infrastructure for Mobile Health and Wellness Monitoring*. Similar systems can be deployed at other institutions to support research and education and to enable students from different disciplines (e.g., computer science/engineering, medical, biomedical, nursing, and health sciences) work together and develop new exciting multidisciplinary health applications and services that may lead to improved quality of life and reduced cost of healthcare.

References

[1] Istepanian, R. S. H., Jovanov, E., and Zhang, Y. Guest editorial introduction to the special section on m-health: Beyond seamless mobility and global wireless health-care connectivity. *IEEE Trans. Inform. Technol. Biomed. 8*, 4, (2004), 405–414.

[2] Gupta, A., Milanesi, C., Cozza, R., and Lu, C. K. Market Share analysis: Mobile phones, worldwide, 2Q13. Gartner, Aug. 2013.

[3] Jahns, R.-G. and Houck, P. Mobile Health Market Report 2013-2017, research2guidance, Mar. 2013; http://www.research2guidance.com/shop/index.php/mhealth-report-2

[4] Quantified Self; http://quantifiedself.com/. [Accessed: 18-Jan-2013].

[5] Milosevic, M., Jovanov, E., and Milenkovic, A. Quantifying Timed-Up-and-Go test: A smartphone implementation. In *Proceedings of the 2013 IEEE International Conference on Body Sensor Networks* (Boston, MA, May 6-9). IEEE, New York, 2013, 302–307.

[6] Milenkovic, A., Milosevic, M., and Jovanov, E. Smartphones for smart wheelchairs. In *Proceedings of the 2013 IEEE International Conference on Body Sensor Networks* (Boston, MA, May 6-9). IEEE, New York, 2013, 404–409.

Biographies

Mladen Milosevic received his Dipl. Ing. Degree in electrical and computer engineering from the University of Belgrade Serbia and his Ph.D. from the University of Alabama in Huntsville in the area of wearable health monitoring. His areas of expertise include ubiquitous health monitoring, smartphone application development, software development, and physiological signal processing.

Aleksandar Milenković is associate professor of electrical and computer engineering at the University of Alabama in Huntsville, where he leads the LaCASA Laboratory (http://www.ece.uah.edu/~milenka). He received the Dipl. Ing., M.S., and Ph.D. degrees in computer engineering and science from the University of Belgrade, Serbia in 1994, 1997, and 1999. His research interests include computer systems architecture, embedded systems, and wearable health monitoring systems. Prior to joining the University of Alabama in Huntsville he held academic positions at the University of Belgrade in Serbia and the Dublin City University in Ireland. He is a senior member of the IEEE, its Computer Society, the ACM, and Eta Kappa Nu.

Emil Jovanov is an associate professor in the Electrical and Computer Engineering Department at the University of Alabama in Huntsville. He received his Dipl. Ing. (1984), M.Sc. (1989), and Ph.D. (1993) from the University of Belgrade. He is recognized as the originator of the concept of wireless body area networks for health monitoring and he is one of the leaders in the field of wearable health monitoring. Dr. Jovanov is a senior member of IEEE, and serves as associate editor of the *IEEE Transactions on Information Technology in Biomedicine* and *IEEE Transactions on Biomedical Circuits and Systems*, and as a member of Editorial Board of *Applied Psychophysiology and Biofeedback*. He is a member of the IEEE Engineering in Medicine and Biology Society (IEEE-EMBS) Technical Committee on Wearable Biomedical Sensors and Systems and a member of the IEEE Medical Technology Policy Committee. Dr. Jovanov has spent more than 25 years in the development and implementation of application specific hardware, software, and systems. His current research interests include ubiquitous and mobile computing, biomedical signal processing, and health monitoring.

Airwriting: Bringing text entry to wearable computers

It may be possible to enable text entry by writing freely in the air, using only the hand as a stylus.

By Christoph Amma and Tanja Schultz

DOI: 10.1145/2540048

Smartphones have emerged as one of the first wearable computing devices to gain widespread usage. Although not a perfect match with the original vision of wearable computing, smartphones are powerful computers available in everybody's pocket. However, the next generation of technologies is around the corner, and such systems follow the idea that we should not have to take an extra device out of our pocket. Instead, they aim to connect us seamlessly with the digital world by presenting information readily at our wrist, as with smart watches, or even directly in our field of view, as is the case with mixed-reality glasses. To interact with such systems, it would be cumbersome to pull our good old smartphone out of a pocket and type a message. Thus, new interfaces are needed to leverage all the possibilities of wearable computing.

One way to proceed is by mimicking human-to-human communication as a prototype for human-machine interaction. Speech, for example, is a very natural form of human communication, but the way people use smartphones indicates text input also plays a crucial role when communicating. People often prefer communicating by text messages, and it is often easier to make small notes by text instead of making a call or recording a spoken memo. The question therefore arises as to how we can realize text entry for wearables. One possibility is to use a wearable keyboard like the Twiddler developed by Prof. Thad Starner's Contextual Computing Group at Georgia Tech [1]. Twiddler is a commercially available device with key buttons that can be worn around the hand. Such a system allows fast and unobtrusive writing, but requires the user to have a device fixed to their palm. Another idea is to project the keyboard and control elements onto any kind of object with a miniature projector worn like a necklace—Pranav Mistry developed such a system at the MIT MediaLab [2].

We expect Airwriting to work best for tasks involving short messages that need to be written while on the move, or for notes that need to be made while performing some other activity.

How about not using a keyboard-based input at all? To provide an alternative to keyboards and other input technologies, we have developed Airwriting, a wearable input system that allows the hand to be used as a stylus. Text can then be entered into a system by performing freehand writing in the air. The hand motion is sensed with inertial sensors that can be integrated into a watch or a bracelet,

Photograph by Volker Steeger


airwriting

meaning neither keyboards nor handheld devices need to be manipulated. The handwriting can be performed against the palm of a person's other hand, mimicking the use of an imaginary notepad (see Figure 1). In this article, we describe some of the challenges we have faced in designing our system and explain the solutions that have made freehand text entry possible with Airwriting.

THE CHALLENGE OF AIRWRITING

The development of a freehand input system faces several challenges. It requires sensing hardware and pattern recognition algorithms to detect and decode the handwritten words from the signals. A sensor capable of measuring motion needs to be attached to the hand or the arm, but this is no longer a big challenge in the presence of Arduino and other low-cost sensor nodes. As there are no special constraints on the sensors, any decent inertial sensor on the market should work for this purpose.

Once the signals are acquired, the handwritten words need to be recognized. As user convenience is a key priority, the system should place as little burden on the user as possible. One undesirable property of technical interfaces is the need for additional control commands besides the actual information necessary for the interaction. In our case, this means that it should not be necessary to require explicit activation, nor deactivation, of the system itself, and the user should not have to make any sort of artificial pauses between characters or words. Furthermore, a user should be able to write immediately as he or she wishes, without having to use a switch or perform some sort of special gesture to start or stop writing.

Once the handwriting is detected, the actual written words need to be recognized. We constrain the system to block capital letters, which are easier for the user to write since no visual feedback of the writing is available. Figure 2 shows an overview of our Airwriting system's components and the processing chain by which gestures are spotted and then recognized as handwriting. Here we will focus on how the system tracks and identifies handwriting—a more detailed description can be found in *Personal and Ubiquitous Computing* for the interested reader [3].

Figure 1: Illustration of the envisioned Airwriting system.

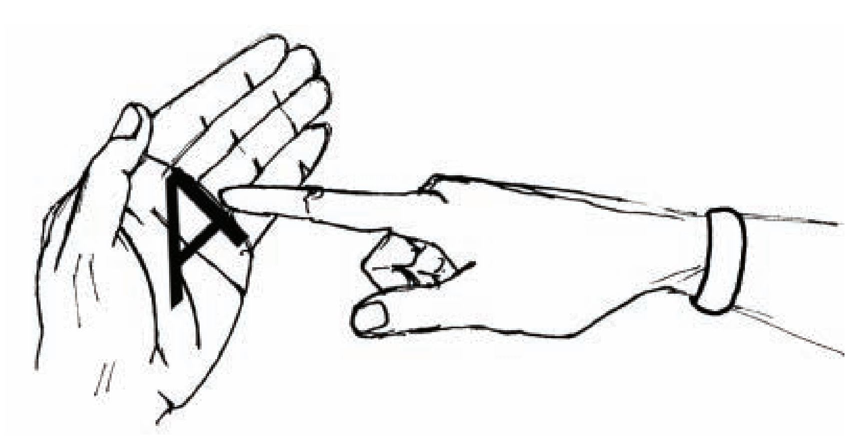

Figure 2: Overview of the Airwriting system's processing chain.

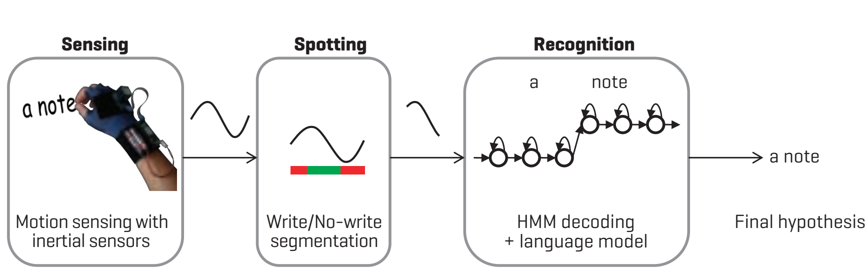

SPOTTING THE GESTURES

Before we can actually recognize what has been written, we need to determine whether or not the user is actually attempting to write something. The act of "not writing" means doing anything else in this case, from cooking to doing sport–just anything. Detection of writing boils down to a binary classification task with one class being handwriting signals and the other being non-handwriting signals. In such problems, the non-handwriting class is often called the 0-class. It contains everything that does not belong to the other class. The 0-class makes such classification problems challenging. The class is not defined by its inherent properties arising from any sort of underlying process, but simply by its distinction from all other classes we are looking at. It follows that it is hard to build a model for this class or to make any conclusion on, for example, the statistical distribution of sensor values. That is because it is impossible to know what, and how often, people do with their hands while they are not writing. Therefore, we need to find characteristic properties of handwriting motions that are rarely seen within the 0-class.

To identify these properties, we gathered sensor data from users performing everyday household activities like cooking, eating, or doing laundry. We used this dataset as an instance of the 0-class. It should be clear the activities we recorded contain only a small part of the 0-class, but it is impossible to record something like a complete dataset for the 0-class. Analyzing the data, we found handwriting movement has characteristics that are not present for the 0-class (for example, a peak in the spectrum at 3 Hz). Figure 3 visualizes the comparison between handwriting and other activities. The shaded columns distinguish patterns of handwriting movement in relation

Figure 3: Example sensor signals during day-to-day activities. The shaded vertical segments contain handwriting, and the non-marked parts belong to the 0-class of non-handwriting movements.

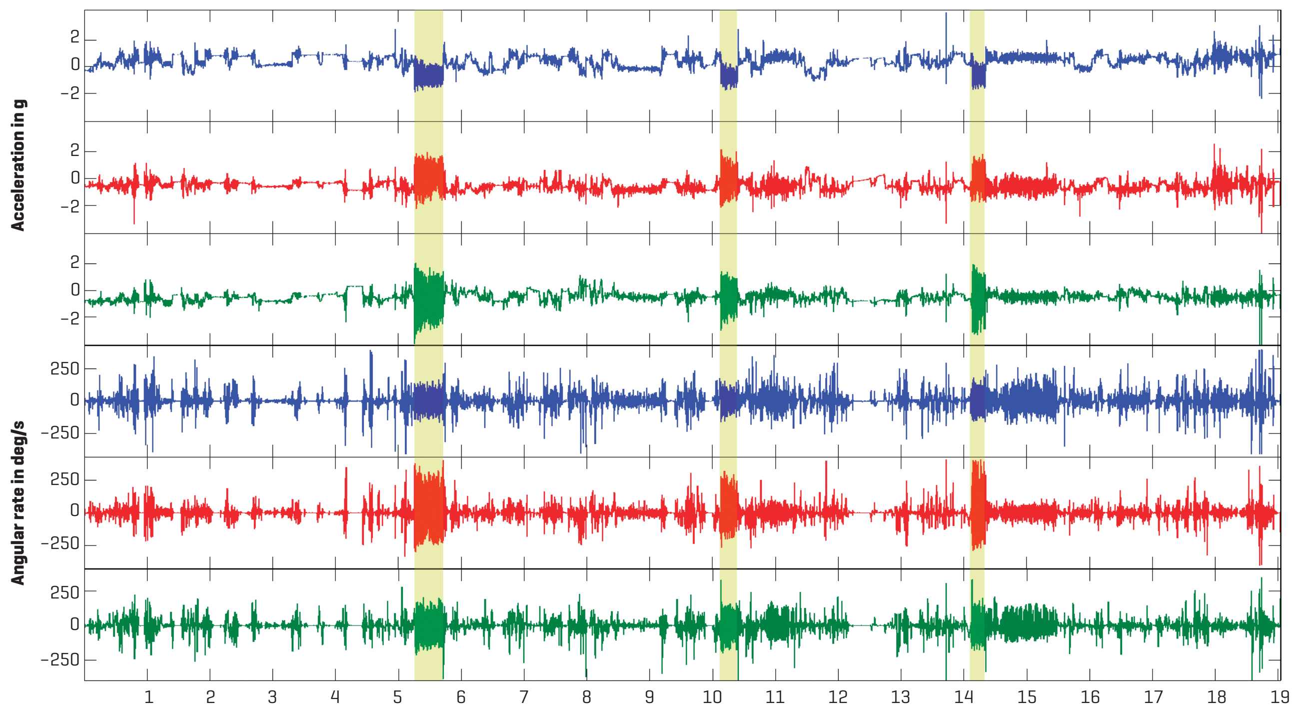

to other actions, with the six rows relating to the six sensor channels we captured.

We used a support vector machine as a classifier and obtained a recall of 99 percent, meaning almost all segments actually containing handwriting were found. However, precision was 26 percent, and thus only about one quarter of the found handwriting segments actually contain handwriting. This seems to be too low for practical usage, but, in this case, high recall is the more important value. If true handwriting segments are dismissed, then they will be lost for the recognition stage. This makes it crucial to identify as many handwriting segments as possible. The low precision value means a lot of segments containing no handwriting are forwarded to the recognition stage. Our analysis showed 98 percent of these segments are so short that they do not lead to a valid recognized word. So, our strategy here is to try to avoid missing any of the possible handwriting segments. This comes at the cost of being somewhat unselective and getting lots of false positives from the spotting stage. We are, however, able to dismiss those wrong segments later based on the recognition results.

RECOGNIZING WHAT IS WRITTEN

Once we have identified segments of the signal that likely contain handwriting, we have to decode the actual written words from the signals. One could think we just have to apply an existing online handwriting recognition system, like the ones that are used for tablets, to identify the text. However, there are two fundamental differences between the signals on which such systems typically operate and the signals we get from inertial sensors. In traditional handwriting recognition systems, we get the actual trajectory of the pen on the surface, i.e. when the pen was really pressed on the surface (pen-down movements). In our case we have neither the trajectory in space nor the distinction between pen-up and pen-down movements. While the trajectory could, in theory, be reconstructed from the acceleration and angular rate signals, the accumulation of sensor errors leads to unusable results within seconds. A second difference is since the pen-up and pen-down information is missing, we are dealing with an entirely unsegmented datastream (as shown in Figure 4), which means we get no information on character or word boundaries from the signal. One way of overcoming this would be to identify pauses, i.e. segments without motion, in the data. However, as stated earlier, we should not put constraints on the user in terms of needing them to make pauses between characters or words. Any such con-

Airwriting saves time and means there is less interruption of a user's primary activity.

Figure 4: Example acceleration signals of the word HAND and illustration of a possible actual trajectory. The red parts of the trajectory indicate what would normally be pen-up movements.

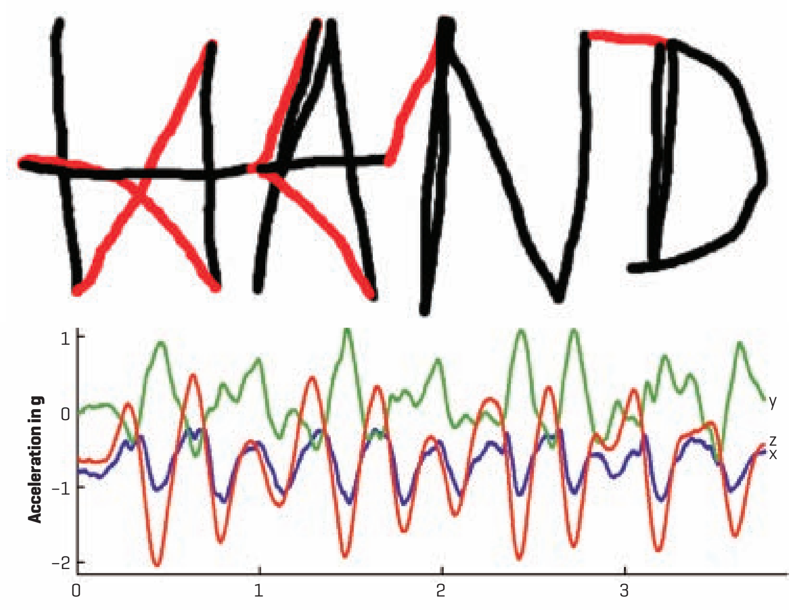

Figure 5: Different variants for writing the letter E observed among different subjects. The grey lines indicate motions that would be pen-up movements when writing on a surface.

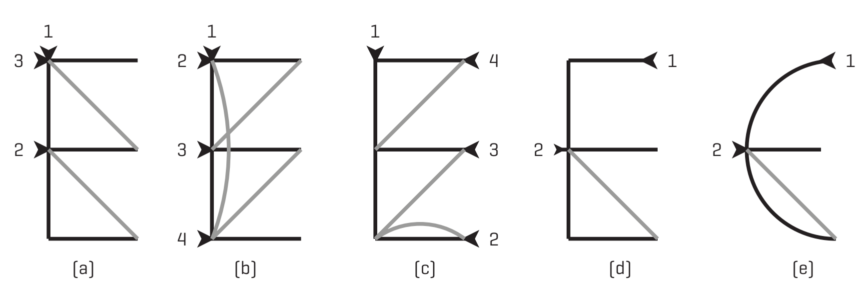

straints would interrupt the fluency of writing and slow down the input.

So, what we are dealing with is effectively a pattern-matching problem. The patterns of individual letters have a characteristic structure in the time domain—the acceleration and angular rate signals look different for an *A* or an *E*, for example. For some letters, the distinction is quite hard. Think of *P* and *D*, *A* and *H*, or *X* and *Y*—these pairs have quite similar motion patterns, making them difficult to distinguish without having the movement trajectory at hand. Additionally, people vary in their way of writing, even for block capital letters (see Figure 5). Since we want people to write words and sentences rather than just characters, we need to recognize sequences of these patterns. However, the boundaries between the individual patterns (i.e. letters) are not known. Hidden-Markov-models (HMMs), which are state-based statistical models, have proven to be well suited to solve such problems [4]. HMMs have been used in speech recognition for years and the challenges in speech recognition are quite similar—a system needs to recognize words as sequences of phonemes without knowing where individual phonemes start or end. HMMs can easily be concatenated to acquire a model for two or more patterns that occur directly after each other. Due to their statistical nature, HMMs can model differences in writing style as well.

In our application scenario, this latter property allows us to build models only for the 26 characters of the alphabet (as we constrain our system to block capital letters). By concatenation of the character models, we can build models for every word we want. This means our system recognizes words, rather than characters, and we can define an arbitrary vocabulary of recognizable words in advance. The current version of our system has a vocabulary of more than 8,000 words. An optimization allows for searching such a high number of different patterns at once. Since all words are formed from the 26 characters of the alphabet, the model for the letter *A* has to be aligned with the signal only once for all words starting with *A*. The same holds for the letters that follow thereafter. One can think of searching through a tree where the nodes succeeding the root node are the 26 possible word beginnings, with each of the leaves then reflecting an individual word.

When it comes to recognizing sequences of words, we can make use of typical properties of language. Not every word sequence has the same probability of occurrence. For example, it is quite likely that we would write the word sequence "do you know" instead of "do you slow." So even if the HMM-based probability for the second sentence was higher, we would choose the first sentence based on statistics of word sequences collected for the writer's language. Such statistics are called language models, and are often obtained by crawling texts on the Web.

The hand motion is sensed with inertial sensors that can be integrated into a watch or a bracelet, meaning neither keyboards nor handheld devices need to be manipulated.

DOES AIRWRITING WORK?

During evaluation studies, our Airwriting system reached a word error rate of 11 percent, meaning, on average, the recognizer makes 11 word errors for every 100 written words. These errors consist of substitutions, insertions, or deletions of words. We also investigated what happens when the system is adapted to the individual user, since even for block capital letters, every user has his or her own style of writing and the sequence of strokes generally differs for some of the letters. As expected, the system's performance improved and the error rate dropped to 3 percent on average. As a given, a wearable interface will probably be used by only one person at a time, therefore the adaptation to the specific user will be a typical use case.

For the evaluation, we asked people to write rather large (10-20 cm) letters as if they were writing on an imaginary blackboard. However, from our own experience with the system, and from the experience gained by showing it at various conferences, we know the writing can get quite small (approximately 3-5 cm) and that the system works well when the user writes at the height of their own hips, as would be the case when using an imaginary notepad (see Figure 1). In terms of writing speed, Airwriting is slower than conventional keyboards. However, if you wanted to enter a relatively short message into a mobile device, you would spend most of the time getting your phone out of your pocket and turning it on before you could write the text. Airwriting saves time and means there is less interruption of a user's primary activity. We expect Airwriting to work best for tasks involving short messages that need to be written while on the move, or for notes that need to be made while performing some other activity.

WHERE TO GO FROM HERE

We already use a variety of different devices depending on the task at hand or the situation in which we find ourselves. Desktop machines are used for tasks where big screens are beneficial; we use tablets when we proofread our text on the couch; and we use smartphones for communicating and performing quick lookups on the go. Int the future, it is likely that we will use a variety of input and output devices to interact with wearable computers. Speech will be one modality, small keyboards will still serve as the best way to input medium sized texts, and gestures can serve as control commands and as a modality to input short texts.

Gestures have several benefits over other input techniques. For example, they can be performed silently without disturbing others. If gestures are small enough, they go unnoticed from bystanders. Last but not least, their sensing is independent from ambient noise, allowing for user input in noisy public places. Therefore, gestures are perfectly suited to complement speech input—they work well when speech won't, and vice versa. Due to the small size of the sensors (a complete inertial measurement unit can be produced by the size of your thumb), they can easily be worn in the form of watches, bracelets, or in the future maybe even in finger rings.

While this article has focused on text entry, the methods we have presented are not restricted to gestures or handwriting. We are currently working toward a more complete interface, using gestures not only for text entry but also for controlling other functions like taking calls, selecting message recipients, or for zooming and scrolling of graphics. Imagine if you could cancel a call by a slight rotation of your wrist, and then send a small note saying that you will call back soon by airwriting it while you are on the go. This is how we envision the future of interaction using wearable technologies.

References

[1] Lyons, K., Starner, T., Plaisted, D., Fusia, J., Lyons, A., Drew, A., Looney, E. W. Twiddler typing: One-handed chording text entry for mobile phones. In *Proceedings of CHI 2004* (Vienna, April 24-29). ACM, New York, 2004, 671–678.

[2] Mistry, P., Maes, P., Chang, L. WUW-wear Ur world: a wearable gestural interface. In *Proceedings of CHI 2009 extended abstracts* (Boston, MA, April 4-9). ACM, New York, 2009, 4111–4116.

[3] Amma, C., Georgi, M., Schultz, T. Airwriting: a wearable handwriting recognition system. *Personal and Ubiquitous Computing* (February 2013), 1–13.

[4] Rabiner, L. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE 77*, 2 (1989), 257-286.

Biographies

Christoph Amma is a doctoral student at the Cognitive Systems Lab at Karlsruhe Institute of Technology. He received his diploma from Karlsruhe Institute of Technology. His interests are in the recognition and interpretation of human motion with body worn sensors. Together with Tanja Schultz, he received the Plux Wireless Biosignals Asward (2010), the best paper award at the International Symposium on Wearable Computers (2012) and a Google Faculty Research Award (2013).

Tanja Schultz is a full professor at the Karlsruhe Institute of Technology (KIT) and the founder of the Cognitive Systems Lab. She has also been a research scientist for seven years at Carnegie Mellon University. She received her Ph.D. from KIT in the area of multilingual speech recognition. Within her team she currently works on human-centered technologies and applications for human-computer-interaction and machine mediated human-to-human communication. In 2012, she received the Alcatel-Lucent research award on technical communication and she is currently president of the International Speech Communication Association (ISCA).

Wearable Brain Computer Interface: Are we there yet?

Brain computer interfaces are still restricted to the domains of health and research, but we understand what needs to be done and are getting closer to making a commercial wearable EEG system.

By Viswam Nathan

DOI: 10.1145/2539268

Brain computer interface (BCI) has been in existence for a long time—UCLA published one of the first BCI reports in 1977—and electroencephalography (EEG) systems exist in many forms. For more than a decade the "P300 paradigm" has enabled the spelling of words on a computer using just our brains; BCIs can also determine which one of multiple flashing objects you're looking at since your brain responds to each of those stimuli differently, known as the SSVEP (steady state visually evoked potentials) task; and using "motor imagery" researchers can tell whether you're thinking about moving to the left or right.

All this sounds wonderfully exciting: P300 could enable texting without using your hands; SSVEP could be used to navigate menus with your brain as each option stimulates your brain differently; while motor imagery has obvious uses in gaming, you don't need to use that analog stick on your controller to turn and look at your opponent during your next game of Halo.

But in order to fully understand why these sci-fi scenarios have not come to pass, it serves to provide some context for the task we are trying to accomplish. For a system to become truly "wearable" for day-to-day use, it needs to be easy to use, provide good performance without undue intervention on the part of the user, and be conveniently small. We will begin with an overview of the hurdles facing EEG systems, and BCIs specifically, before they can meet the above definition of a wearable system.

Brain computer interface in its current state is still too slow, cumbersome, unreliable, and impractical for casual day-to-day use.

CHARACTERISTICS OF EEG

First, we need to understand the type of signals we are dealing with to provide some context. EEG is the measurement of the brain's electrical activity. Electrodes placed on the scalp are designed to pick up this electrical activity. When the brain initiates a task, the stimulus to complete the required action is communicated through the neural network.

The neural network is composed of neurons, which are electrically excitable cells. One neuron can electrically excite its neighbors thus building a chain of electrical signaling. When the electrical excitation is high enough, a voltage difference called the "action potential" is achieved and this traverses the neurons in the network. The electrodes placed on the

Photograph by Janne Moren

THE 5 MAIN HURDLES OF WEARABLE BCI SYSTEMS

1. Wet electrodes are impractical, but dry electrodes have a high impedance contact.
2. Low SNR for BCI tasks means slow processing time.
3. Low SNR also dictates the necessity for multiple electrodes placed around the head.
4. Training time is required for some paradigms, such as P300, to calibrate the system.
5. EEG is still very susceptible to motion artifacts and most currently implemented BCI tasks require you to stationary.

head are designed to read this wave of electrical activity going across the scalp. As you might think, different tasks produce different patterns in different parts of the scalp. We are still not at the level of distinguishing every possible thought or process, but there are certain readily identifiable patterns. For example, if you close your eyes there is a steady rhythm in the EEG in the frequency range 8–12 Hz known as the alpha. There is also more activity in the left side of your brain when you think about moving your right hand and vice versa.

So just how strong are these signals? The generated action potential can be as high as 100 millivolts (mV). You may think this small, but in reality it is no issue for the equipment that we have today. However, this is the signal strength right at the neural level and not at the scalp. The signal gets dissipated significantly by the time it reaches your skull. There are ways to surgically introduce electrodes inside the scalp to read the signal directly, but unless it's for purely medical purposes nobody would be happy going through a procedure that is the very definition of invasive. So we need to measure at the scalp surface. Unfortunately, skin is not a great conductor and hair is even worse.

If an electrode manages to penetrate through the hair and make contact with the scalp, the signal being read is on the order of a few microvolts. For comparison, the electrocardiogram (ECG) pulse on the surface of your skin can be a few millivolts. Moreover, any amount of movement on the electrode can cause motion artifacts that absolutely swamp the EEG signal. Furthermore, this signal in microvolts represents all brain activity and not just the signals of interest. In other words, there is a lot of background activity that is irrelevant to the BCI task at hand. Although the brain needs to do thousands of other things at the same time as your BCI task—to keep you alive and functional—from a system development point of view this background activity is noise and reduces the signal-to-noise ratio (SNR). An example of this is the P300 task I mentioned earlier.

CHALLENGES OF BCI: THE P300 SPELLER

P300 has already been in use for several years to help people spell using their brain. The user is presented with a matrix of letters and numbers on a screen, and is asked to focus attention on the specific character he/she wishes to spell.

While the user is focusing on one character in the matrix, random rows and columns of the matrix start flashing on and off for brief periods. Occasionally, the particular character being focused on by the user will also flash brightly for half a second or so, and this sudden flash triggers a known "peak"-like response in the user's EEG waves 300 milliseconds after the fact: Hence the name P300. The flashing must appear random to the user in order to generate a "surprise" effect every time their object of focus illuminates. A P300 will not be elicited if the user is expecting the stimulus.

Since the flashing is so brief, in theory, it should take no time at all to successively flash all rows and columns in the matrix once each, then determine which row and which column produced the P300 response, and pinpoint which character was being observed. However, in reality the evoked P300 is so small compared to the rest of the EEG activity that we need to average the data over several trials in order to predict the correct letter with any confidence. In our lab, for example, the P300 speller we built requires at least 15 flashes each for rows and columns and it takes about 60 seconds to spell a single character. SNR again is the major burden since the required signal is so miniscule that we need to see it multiple times to be confident that it is there.

Apart from all this, every P300 system I'm aware of requires some training before the user can start spelling. As you might expect, different people have different brain signals and respond differently to the same stimuli. Essentially you have to teach the system what exactly your P300 looks like. Fair enough you say, we can put up with a little bit of training time for the system. But it doesn't stop there: If you want to use the system again at another time you'll have to train it again. EEG can be quite finicky and what you might think are small differences (electrodes are not placed the same way as last time, your cap is slightly shifted, you have a little more sweat on your head, etc.) can all cause significant differences in your signal. Essentially meaning the system has to throw away what it learned about you in the past and start over. Training with the system in our lab takes about 30 minutes if everything goes smoothly.

It is still a rewarding experience using the system the first couple of times—"I'm actually spelling words using just my brain!"—but ultimately in its current state, even the best case scenario involves a lot of hassle and is painfully slow.

The other major hindrance to this system becoming wearable is the number of electrodes that need to be placed on the head. In our lab we use eight electrodes placed over the top, sides, and back of the head and even this is a low number compared to other work in this area. The reason for this is to again build some redundancy since not every single electrode will pick up the P300 cleanly. In a sense, the lack of quality of the data is being compensated for by increasing the quantity. Unfortunately this means putting up with the discomfort of several electrodes poking through your hair (not to mention the unseemly appearance). Undoubtedly the P300 in its current form is invaluable to patients suffering from diseases, such as ALS (or Lou Gehrig's disease), who are paralyzed. They have gained a mode of communication where previously there was none. But in the context of a wearable system to be used in day-to-day life, there's still a long way to go.

For a system to become truly "wearable" for day-to-day use, it needs to be easy to use, provide good performance without undue intervention on the part of the user, and be conveniently small.

EEG ELECTRODE DESIGN

Research into motor imagery must overcome similar challenges as the P300. Motor imagery tries to infer from your brain activity the direction in which you're thinking about moving. More sophisticated algorithms can also tell if you're thinking about moving your hand or foot. This could have a number of exciting uses in helping people use prosthetic limbs or, as I mentioned before, in gaming. However in its current state it requires significant training time, has a fairly low SNR, and requires multiple electrodes to properly map the brain activity.

This all means that careful design of the electrodes themselves is vital. The most popular design for EEG systems, especially in the medical field, adopts "wet" electrodes. These electrodes basically enclose a conductive silver-silver chloride (Ag-AgCl) gel solution. With proper scalp preparation a suitable low impedance contact can be made. The signal quality is indisputable for this design and they are better for medical diagnoses than any of the alternatives simply because they are less susceptible to external noise sources by virtue of the low impedance contact. However there are drawbacks to this system:

1. Your scalp needs to be abraded so the gel in the EEG electrodes can make good contact. (This is a painful process and unfortunately it must be repeated each time you use the system, as your skin will regrow. Scalp abrasion also carries the risk of infection.)
2. Another person needs to be present to inject the gel into the cap.
3. The gel itself can get quite messy and cause irritation if left on your head for long periods.
4. It requires a significant amount of preparation time before any readings can be taken.
5. The acquisition system, along with all the wiring, can be quite cumbersome.

As you can see this does not fulfill any of the criteria for a wearable system and suddenly playing Halo with your analog stick doesn't seem so bad anymore. That's also why EEG acquisition systems have been limited to medical and diagnostic purposes so far. It's the only scenario where the possible benefits outweigh the annoyances from the patient's perspective.

Recent research has gone into the development of dry EEG systems to overcome some of these drawbacks. Dry electrodes, as the name suggests, do not require the use of a conductive gel and most designs consist of some kind of metal surface placed directly on the head. No more scalp abrasion, no more messy gel, and a drastic reduction in preparation time. The catch is the aforementioned SNR takes a hit. The reason conductive gel is used in the first place is to reduce the impedance of contact, so that we can see a stronger signal relative to the surrounding noise. Improvements in amplifier and ADC technology have meant that we can feasibly use dry electrodes, but the impedance faced now is a few orders of magnitude higher. This basically means the noise floor is much higher and it doesn't take much to completely disrupt the overall signal and compromise any BCI task. Dry electrodes are also much more vulnerable to motion artifacts, which disrupt the signal.

Regardless, the everyday user is not going to purchase a wet electrode system just because of its higher SNR. Dry electrodes are the way forward, and so we need to find other ways to compensate for the low signal level.

It's not all doom and gloom; let us now look at the bright side of things.

BCI ACCOMPLISHMENTS

We've come a long way already and there are certainly a few exciting BCI applications within our reach. There are already commercial BCIs with dry electrodes that are comfortable to wear and use, like Emotiv for example. They are limited to games based on simple tasks like detecting alpha waves—the tell-tale 8-12 Hz EEG rhythms that are elicited when a person is relaxed—but it's certainly a good first step.

More exciting are the paradigms based on SSVEPs. When you look at a periodically flashing object, let's say an LED flashing at a frequency of 7 Hz, the brain signals originating in the visual cortex at the back of your head assume a discernible pattern. Funnily enough, the dominant frequency of this pattern exactly matches the frequency of the flashing stimulus you are currently looking at. So if we put three LEDs in front of you, each flashing at a different frequency, we would be able to tell which one you are looking at just based on your EEG signals. This has already been implemented in several fun applications: In our lab we have a rock-paper-scissors game where instead of performing the gesture using your hand, you just look at the appropriate flashing object on the screen and the computer knows which one of the three you will chose to play. There's already been a lot of work extending this beyond fun and games to useful tasks such as dialing a number on your phone by just looking at your keypad—where each number is flashing at a different frequency.

The great thing about SSVEP is it requires no training time. Of course, some people may respond to certain frequencies better than others, but, on the whole, the signal processing is fairly robust and simple. The signals are also fairly localized at the back of your head and we do not need to map as much brain activity as other tasks. Most SSVEP BCI implementations still use about eight electrodes, but we've had some success in our lab using just a single electrode.

Motor imagery tries to infer from your brain activity the direction in which you're thinking about moving. More sophisticated algorithms can also tell if you're thinking about moving your hand or foot.

ADDRESSING WEARABLE BCI CHALLENGES AND LOOKING AHEAD

To conclude, I'll briefly summarize my own research in making a wearable BCI system, with focus on the design and development of reconfigurable BCI systems. Specifically, I want to make the hardware a little smarter to both increase SNR as well as reduce the setup time and hassle of using an EEG system. Currently, even dry electrodes are difficult to use and setup may require assistance from another person. Each electrode needs to be properly pressed down and twisted through hair to make enough contact with the scalp to ensure the signals are not too noisy.

I am working on a method to continually monitor the contact quality of these electrodes with the scalp and respond to changing conditions in real time. In other words, the subject would just have to put the EEG cap on reasonably well and the hardware will automatically sense the conditions and reconfigure to achieve the optimal contact. Moreover, this information on contact quality can also be leveraged in the back-end signal processing. For example, when processing a motor imagery task it would be useful to know if the electrodes on one side of the head are not placed as well as on the other side so this can be accounted for in the signal processing. Or take the P300 paradigm discussed earlier, what if I could tell the system exactly how much worse or better the contact of the electrodes is from the training data collected a week ago? How this would affect the current data? It could potentially rid us of the need for multiple training sessions, which is certainly an inviting prospect. The monitoring of contact quality could also pave the way for signal processing techniques to reject motion artifacts. After all, motion of the electrode affects the impedance of contact with the scalp, but does not affect the underlying EEG activity. Recent research has already shown these changes in impedance correlate well with the type of motion being performed; so the prospects are certainly good for the ability to adaptively reject the motion artifacts by leveraging continuous monitoring of the electrode contact.

Brain computer interface in its current state is still too slow, cumbersome, unreliable, and impractical for casual day-to-day use. However the problems, and most of the solutions, are well defined. Several researchers around the world are steadily making dents in these obstacles. Recently, there have been many advancements in the hardware and signal processing to enhance SNR, reduce the training time, increase robustness to motion artifacts, improve dry electrode quality, and decrease the number of electrodes necessary to perform a BCI task well. For the layperson, EEG and brain computer interface have just been exciting buzzwords, but we're well on the way to making wearable BCI a reality.

Biography

Viswam Nathan is currently pursuing his Ph.D.in computer engineering at the Embedded Systems and Signal Processing Lab at the University of Texas at Dallas. His research interests include design of dry contact EEG electrodes, development of a reconfigurable brain computer interface, as well as techniques for motion artifact rejection for both ECG and EEG.

PROFILE DEPARTMENT EDITOR, ADRIAN SCOICĂ

Ori Inbar
Making Augmented Reality a Reality

DOI: 10.1145/2544054

When I sat down to interview Ori Inbar on a late Wednesday evening, I already knew he was one of the most influential figures in the world of augmented reality. Inbar is CEO and co-founder of Ogmento (the first venture-based company developing augmented reality games), CEO and co-founder of a non-profit organization called AugmentedReality.ORG (which founded the Augmented World Expo, the largest international conference of its kind), and the president of the Augmented Reality Consortium. What I didn't know was that he is a passionate visionary with a plan, whose personal story and drive make his career accomplishments even more inspiring.

Inbar graduated from Tel Aviv University with a double major in computer science and cinema, and then joined a startup in the early '90s, where he developed multimedia and business software. The startup was later acquired by ASP, and he stayed on to develop a new platform called NetWeaver until 2007, when he realized he wanted to pursue something bigger and more meaningful after having already worked for startups and large corporations.

Since finding one's true calling is a nearly-universal human dilemma, I wanted to know how he discovered augmented reality. To my surprise, he confessed: "It's one of those cases where I don't know if I discovered it, or it discovered me. I came home one day and realized my kids were always stuck in front of the screen. It seemed this is how we communicate in the 21st century, but it also felt like we were missing a lot of it in the real world, so I was looking for a way to extract what attracts kids to video games, and bring it into the real world.

That's when I realized augmented reality has been around for decades, but it's been hidden in the labs. My mission became to find a way to bring it to the masses."

Although it is tempting to only think of augmented reality in terms of video games and entertainment, Inbar explained there already exist applications over a wide range of domains, from education to warehouse picking: "You can imagine glasses, or a mobile device, which points you to where you need to pick a certain item in a warehouse, and once you do it, it automatically updates the inventor." According to Inbar, large corporations are very interested in that sort of application. However, one of his favorite applications is AR Pool. It is a pool table that tracks the balls and the position of the cue, and shows you how to send the ball into the pocket every time, thus demonstrating how augmented reality can benefit education by allowing novices to instantly master skills.

It is no wonder that with so much untapped potential in the field, the mission of AugmentedReality.org is to help advance the field and the whole ecosystem around it. To explain what the latter implies, Inbar enlisted three action verbs that describe his work: connecting, educating, and hatching.

"Connecting is about bringing people together, and we do it by local meetups all over the world, as well as the Augmented World Expo, which is now in its fifth year," he said. The event allows people to form partnerships and to network in order to find the talent and technology they need. Educating is, of course, the ongoing effort of helping companies figure out how to leverage these new technologies so that they are successful, and not perceived as gimmicky. But perhaps most importantly, hatching is about helping startups with expertise, potential funding sources, and business insights through mentoring.

Together, he claimed, the three goals of his work aim to create an ecosystem encompassing both startups and large industry players such as Google, Samsung, or Intel, which help drive the field further. To illustrate his point, Inbar explained how Google's latest gadget helped boost an entire startup environment: "Glass is one of the best things that happened to augmented reality in the past few years, just because of the sheer marketing power of Google. Although Glass is more of a notification environment, and doesn't truly fit the definition of augmented reality because it doesn't overlay graphics on the world, it's bringing a lot of attention to the field, and now all of a sudden you have a stage set for many other companies that have also been developing augmented reality glasses for years; these glasses are perhaps in many ways more advanced products than Google Glass, but less known."

According to Ori Inbar, the future holds a whole new and exciting way of interacting with the world. He cited analyst Tomi Ahornen, who described augmented reality as the eighth mass medium and—based on the current adoption rates—predicted one billion users by 2020. For a quick taste of the future, Inbar highly recommended reading *Rainbow's End* by Vernor Vinge and *Daemon* by Daniel Suarez, two books that helped him reframe science fiction as an attainable objective.

"Personally, I wake up and go to bed, and sometimes even dream about augmented reality, so it's definitely part of everything I do," he proudly related.

end

The David R. Cheriton School of Computer Science at the University of Waterloo.

LABZ

Cryptography, Security and Privacy (CrySP) Research Group

Waterloo, Canada

My pursuit of a research career in information privacy brought me to the Cryptography, Security and Privacy (CrySP) group at the David R. Cheriton School of Computer Science, University of Waterloo. The group offers a fertile ground for research on a diverse range of cryptography, security, and privacy topics. These topics can be broadly divided into the following areas: cryptographic efficient algorithms and distributed protocols (Doug Stinson); security and privacy enhancing technologies (Ian Goldberg) focused on usefulness, effectiveness and usability of cryptographic and security systems; location based privacy (Urs Hengartner); privacy for social networking and online voting (Urs Hengartner); and information systems assurance and management (Ian McKillop) for financial and health sectors. There is the core team of the principal researchers, as well various other faculty members. As a result, the CrySP group attracts a fair number of students. Currently, the group supports 15 active students and it has a long list of successful alumni (31 since 2007).

Jalaj Upadhyay is a graduate student at University of Waterloo under the supervision of Prof. Douglas R. Stinson. As a member of CrySP he is involved in two projects. He shared the following about his experience in the CrySP group: "One component of my research is understanding the role of cryptography in the domain of cloud computing. My research focus in this domain is in understanding the security requirements, which capture the real-world requirements of a user that either stores a file on a cloud or performs some computation on its stored data. A continuation of this research direction is construction of efficient protocols that satisfies the security requirements." At CrySP we are free to explore more than one research area. In Jalaj's case, his second project is focused on differential privacy, a very robust guarantee of privacy on social networks. He explains, "I am interested in investigating whether we can construct efficient differentially private mechanisms that answers certain set of queries on social network without leaking any information about an individual."

Numerous research projects at CrySP have had a positive impact on many useful real-world privacy and security related applications. The Tor project is a shining example of such an effort, where the research work of many individuals contributes toward enhancing various privacy and security concerns. Tor provides an open source implementation of the onion routing protocol, as well as an open network for online individual anonymity. Off-the-Record Messaging (OTR), which enables secure and private instant messaging (IM) over existing IM networks, is another CrySP success story. OTR provides encryption, authentication, deniability, and perfect forward secrecy out-of-the box for a large number of IM clients. Percy++, an open source implementation of private information retrieval (PIR) protocols, is another project that benefited greatly from the research done by the members of the CrySP group. FaceCloak, for protecting user privacy on social networking sites, is made available under an open source license, along with a default Firefox plugin based implementation. (There are numerous other projects that can be re-

At the University of Cambridge, the Trojan Room coffee maker was the first Internet-connected appliance set up with a webcam to show how much coffee was left in the pot.

viewed at the CrySP website.)

CrySP researchers are also involved in the theoretical aspects of security and privacy. Another area of research is the application of combinatorial objects, like design theory, in the construction of efficient cryptographic protocols kept secure against an adversary that has unlimited computational power. Efficient batch zero knowledge proof is a research project that allows a party to prove to another party convincingly the validity of a statement without giving out any knowledge of how the statement is true.

As a Ph.D. candidate under the supervision of Dr. McKillop, my work focuses on the expression and the (real-time) enforcement of individual consent across heterogeneous information systems. I am interested in creating consent-centric access control models that can provide automated reasoning for access control decisions across different administrative domains. These consent-aware models are used as building blocks for designing distributed multiparty access control protocols to ensure enforcement of consent directives for real-time exchange of personal information.

In addition to the standard cryptographic techniques, my work relies on artificial intelligence primitives such as structured knowledge representation, logic-based reasoning, and decision under uncertainty to provide the desired privacy guarantees. In order to validate my research, I have chosen the consent specific use cases from the complex healthcare domain. It is my hope that my work will empower patients to express rich consent preferences, which can be reasoned about by automated systems.

Biography

Atif Khan is a Ph.D candidate under the supervision of Ian McKillop at the David Cheriton School of Computer Science, University of Waterloo. He is interested in solving the patient consent challenges for medical information management.

BACK

Robotic Vacuums

There is no question that our homes, places of work, and world in general are becoming more inundated with technology designed to make our lives more convenient with each passing year. Some technologies stick and some do not. In recent years our homes have seen improvements in networking, intelligent thermostats, integration with modern electrical grid technologies, as well as domestic robots. One technology that has enjoyed some remarkable and rapid success is the robotic vacuum cleaner.

The first robotic vacuum to reach the commercial market was the Trilobyte by Electrolux. It included features that were impressive at the time, many of which are common in more recent models. For example, when the battery charge was low, it automatically navigated back to its base for recharging before returning to the same spot to finish the job. It used ultrasound for collision avoidance as opposed to the bump sensors more commonly found today, and even mapped its workspace for precise cleaning and navigation to the base station. The first generation Roomba by iRobot followed the Trilobyte to market in 2002 as a cheaper alternative. As for features, the Roomba was comparable in most respects, with the primary difference being the algorithm it used for cleaning. The Roomba didn't map the room, but followed a set of simple movement patterns as well as some random walk behavior. This method was perhaps less sophisticated, but proved to be effective. The Roomba line operates in much the same way to this day.

There is still plenty of room for making these devices truly autonomous, but demand will undoubtedly help to drive development. Given how quickly they have come in the last 10 years, it is interesting to imagine how far they may yet go in the next 10. —*Finn Kuusisto*

Trilobyte (left) by Patrik Tschudin. Roomba (not a 790, but close) by Tibor Antalóczy.

	Trilobyte	**Roomba 790**
Manufacturer	Electrolux	iRobot
Release Year	2001	2012
Price	$1,500	$700
Automatic Charging	Returns automatically to charging station	Returns automatically to charging station
Charge Time	2 hours	3 hours
Battery Life	1 hour	1 hour
Sensors	Ultrasound, magnetic	Bump, IR, cliff
Algorithm	Maps the room	Follows simple patterns as well as random-walk behavior

HELLO WORLD

On Constructing the Tree of Life

BY MARINKA ZITNIK

Evidence today suggests both living and extinct organisms are genetically related. These genetic relationships can be represented by an evolutionary tree called the "tree of life," dating back to times of Charles Darwin in the early 19th century. The modern development of this idea are phylogenetic trees—diagrams of relatedness between organisms, species, or genes—that show a history of descent from common ancestry. As more and more life sciences data are freely available in public databases, some of the analyses that would have been performed in well-equipped research laboratories just few years ago are nowadays accessible to any interested individual with a commodity computer. Such a shift was only possible due to unprecedented technological and theoretical advancements across a broad spectrum of science and technology. Herein we describe a simple, but complete, pipeline that includes acquiring genetic data from an online biological data repository and constructing a phylogenetic tree, including its visualization and interpretation.

Obtaining Data

Our aim is to generate a phylogenetic tree of selected species (see Definition 2) from their biological sequences. Specifically, our analysis will base on protein apoptotic protease activating factor 1, also known as APAF1, whose homolog has been found in all currently sequenced animal genomes [1]. A homolog gene is related to another gene by descent from a common ancestral DNA sequence. In this tutorial we use Biopython 1.6 to fetch protein sequences from a biological data repository, ClustalW2 tool to align the sequences, and NetworkX 1.8 for visualizing phylogenetic trees. It is a common task in bioinformatics to extract information from biological databases. Definition 1 shows how to retrieve the sequence data from the Entrez database and save them to a local file using Biopython library.

Creating the Phylogenetic Tree

Most existing approaches for phylogenetic inference can be divided into two groups. Algorithms in the first group compute a matrix of distances between each pair of biological sequences and transform this matrix into a tree. Other techniques find a tree that best explains the observed sequences under a selected evolutionary model by evaluating the fitness of different tree topologies. In this tutorial we rely on the approach from the first group in which we construct a phylogenetic tree by aligning protein sequences.

Once we find the correct sequences, we align them to see how similar they are [2]. We use the ClustalW2 tool to do such an alignment. This is a general-purpose global multiple sequence alignment tool that provides biologically meaningful alignment of divergent sequences. It calculates the best match for the selected sequences across their entire length and lines them up, so that the similarities, variations, and identities can be easily observed.

ClustalW2 generates an alignment and guide tree files with names based on the input FASTA file (see Definition 3), in our case "apaf1.aln" and "apaf1.dnd", respectively. The latter is just a standard Newick tree file, which can be parsed with Phylo module in Biopython. If you installed the prerequisites you can export the tree object to a NetworkX graph, use Graphviz to lay out the nodes and display it with matplotlib (see Definition 4). Otherwise, it is possible to create an ASCII-art dendrogram with the `draw_ascii` function in Phylo or by using other tools such as iTOL [3]. The `rooted` attribute of Phylo tree object creates a head on each branch to

Definition 1: A Python script to download APF1 protein sequences of selected species from NCBI GenBank database and to save the sequences to a FASTA formatted text file.

```
from Bio import Entrez
Entrez.email = 'me@uni.edu'
gb_id = 'NP_863651.1, XP_003313928.1, XP_001087067.1, XP_003432082.1,' \
    'NP_001178436.1, NP_001036023.1, NP_076469.1, XP_416167.3, NP_571683.1,' \
    'AFN55258.1, ABJ16405.1, AEX93473.1, XP_005143187.1, XP_005039610.1, ' \
    'XP_005024582.1, XP_003476187.1, XP_004787271.1, XP_004650259.1, ' \
    'XP_004623872.1, NP_001085834.1, XP_003221150.1, XP_004269550.1,' \
    'XP_004381311.1, XP_003907055.1'
h = Entrez.efetch(db='nucleotide', id=gb_id, rettype="fasta", retmode="text")
f = open('apaf1.fasta', 'w')
f.write(h.read())
f.close()
```

Definition 2: Species used for creating a phylogentic tree. Dictionary defines mapping between Entrez accession numbers and species names.

```
name2org = {
    None: '',
    'gi|32483359|ref|NP_863651.1|': 'H. sapiens',
    'gi|332840131|ref|XP_003313928.1|': 'P. troglodytes',
    'gi|109098377|ref|XP_001087067.1|': 'M. mulatta',
    'gi|345781094|ref|XP_003432082.1|': 'C. lupus',
    'gi|330864777|ref|NP_001178436.1|': 'B. taurus',
    'gi|110347465|ref|NP_001036023.1|': 'M. musculus',
    'gi|13027436|ref|NP_076469.1|': 'R. norvegicus',
    'gi|363727703|ref|XP_416167.3|': 'G. gallus',
    'gi|18858279|ref|NP_571683.1|': 'D. rerio',
    'gi|395395167|gb|AFN55258.1|': 'C. orientalis',
    'gi|115607117|gb|ABJ16405.1|': 'F. catus',
    'gi|372471328|gb|AEX93473.1|': 'S. mediterranea',
    'gi|527249745|ref|XP_005143187.1|': 'M. undulatus',
    'gi|524983223|ref|XP_005039610.1|': 'F. albicollis',
    'gi|514766988|ref|XP_005024582.1|': 'A. platyrhynchos',
    'gi|348580841|ref|XP_003476187.1|': 'C. porcellus',
    'gi|511936741|ref|XP_004787271.1|': 'M. putorius',
    'gi|507532372|ref|XP_004650259.1|': 'J. jaculus',
    'gi|507616819|ref|XP_004623872.1|': 'O. degus',
    'gi|148231147|ref|NP_001085834.1|': 'X. laevis',
    'gi|327272756|ref|XP_003221150.1|': 'A. carolinensis',
    'gi|466007654|ref|XP_004269550.1|': 'O. orca',
    'gi|471394804|ref|XP_004381311.1|': 'T. manatus',
    'gi|402887344|ref|XP_003907055.1|': 'P. anubis'
}
```

Definition 3: Aligning protein sequences.

```
from Bio.Align.Applications import ClustalwCommandline
cline = ClustalwCommandline('clustalw2', infile='apaf1.fasta')
cline()
```

Definition 4: Visualizing a phylogenetic tree.

```
import networkx as nx
import matplotlib.pyplot as plt
from Bio import Phylo

tree = Phylo.read('apaf1.dnd', 'newick')
tree.rooted = True
G = Phylo.to_networkx(tree)

G1 = nx.convert_node_labels_to_integers(G, first_label=1)
pos = nx.graphviz_layout(G1, prog='neato', args='')
posn = {n:pos[i+1] for i, n in enumerate(G.nodes())}
labels = {n:name2org[n.name] for n in G.nodes()}

nx.draw(G, posn, node_size=0, labels=labels)
plt.show(G)
```

indicate direction of its edge. The `prog` argument specifies the Graphviz layout engine. We utilize the neato program, which produces useful visualizations of moderately-sized trees.

Reading a Phylogenetic Tree

A tree structure is helpful in tracking biological diversity at all levels, whether we consider the most diverse branches of the tree of life or recently diverged lineages. A typical phylogenetic tree is a rooted tree where the root corresponds to the common ancestor of all species included in a tree and the tips of a tree correspond to individual organisms, to species, or to sets of species. A general term for the tip of a phylogenetic tree is a taxon (i.e. a leaf of a tree) and the lines represent evolutionary lineages. The branching points within a tree denote speciation events and are called nodes. A node corresponds to the last common ancestor of the two lineages descended from that node. Part of a phylogeny that includes an ancestor and all of its descendants is called a clade. This group of taxa has the property of monophyly (from the Greek for "single clan") and is referred to as a monophyletic group. It can be easily identified visually: It is a piece of a larger tree that can be cut away from the root with a single cut. For instance, *H. sapiens*, *P. troglodytes*, *P. anubis* and *M. mulatta* taxa in Figure 1 form a monophyletic group. Clades define biologically interesting parts of phylogenies because all members of a clade share some part of history, which consists of common ancestry on the internal branch that attaches a clade to the rest of a tree. For instance, if we are told turtles share a more recent common ancestor with birds than with snakes and lizards and we are given a clade of turtles, then we can conclude all turtle species share a more recent common ancestor with birds than with snakes or lizards.

Phylogenetic trees usually depict only the branching history of ancestry.

Figure 1: A phylogenetic tree showing evolutionary relationships between 24 species including human based on alignment of their APAF1 protein sequences.

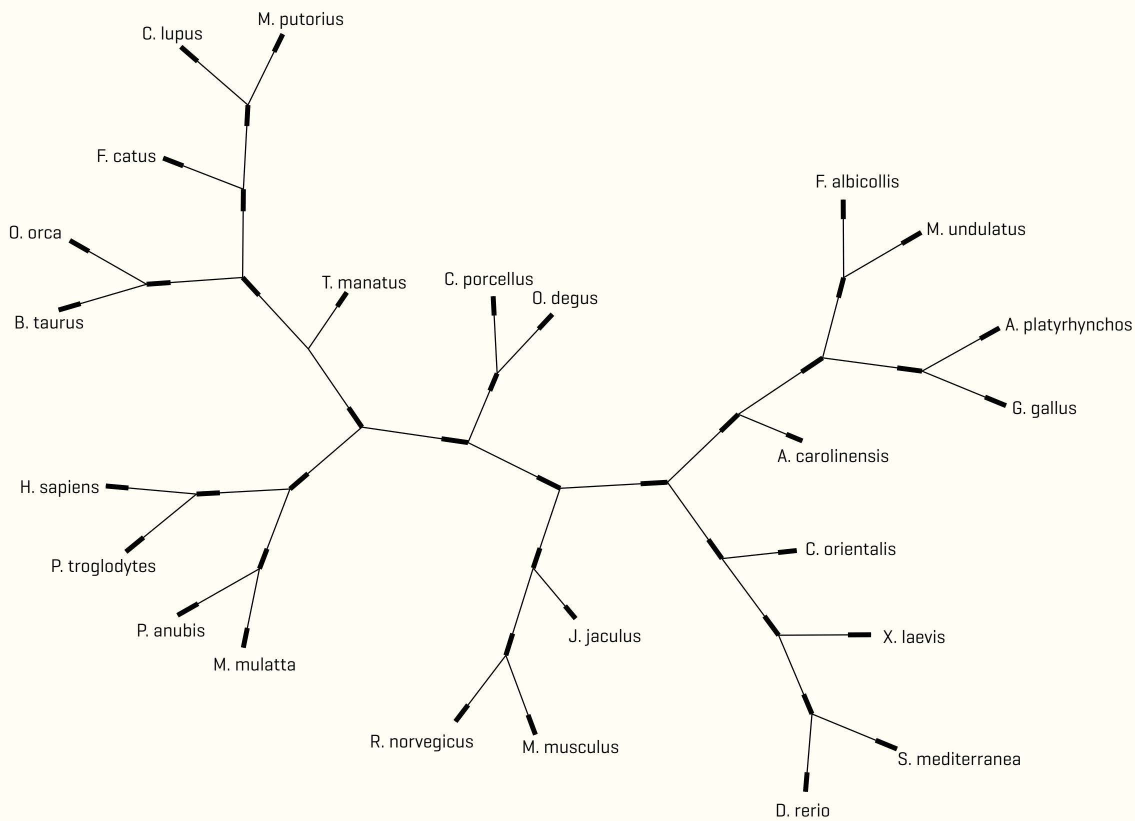

When reading a phylogenetic tree one is interested in the pattern of branching and the overall topology of a tree and not in the lengths of branches. A tree contains the same information about evolutionary descent regardless of its branch lengths unless stated otherwise. We sometimes draw tree branches such that their lengths are relevant and represent the amount of evolution in genetic material or estimated duration of lineages. When this is not true we should avoid reading any temporal information from a tree. For instance, Figure 1 may suggest a speciation event leading to tips *C. lupus* and *M. putorius* occurred after the event that separated tip *O. orca* from *B. taurus*. However, in reality (*C. lupus*, *M. putorius*)-node could have happened either before or after (*O. orca*, *B. taurus*)-node. Evolutionary implications are independent of tree orientation in space and position or shape of its branches. The basic rule is that two trees depict the same evolutionary history if you can transform one tree into another by twisting, bending or rotating branches without cutting them.

Conclusion

The use of phylogeny is expanding to many areas of biological science and represents an essential tool for organizing the knowledge of evolutionary history. This tutorial provides a basic example of how developments in sequencing technologies and genome analysis methods can be used to investigate biology. Current research goes beyond theory to provide new insights into our health and disease that will shape our everyday lives. We are at the outset of a new era, where patients and doctors will have access to whole genome sequences to customize medical treatment and tailor medical care to each individual [4]. There are already many examples of personalized medicine in current practice, such as reducing the incidence of adverse drug effects by checking for susceptible genotypes and developing patient-specific drug dosing algorithms [5].

References

[1] Cecconi, F. et al. Apaf1 (CED-4 homolog) regulates programmed cell death in mammalian development. *Cell* 94, 6 (1998), 727-737.

[2] Chenna, R. et al. Multiple sequence alignment with the clustal series of programs. *Nucleic Acids Research* 31, 13 (2003), 3497-3500.

[3] Letunic, I. and Bork, P. Interactive Tree Of Life (iTOL): An online tool for phylogenetic tree display and annotation. *Bioinformatics* 23, 1(2007), 127—128.

[4] Fernal, G.H- et al. Bioinformatics challenges for personalized medicine. *Bioinformatics* 27, 13 (2011), 1741-1748.

[5] Sagrieya, H. et al. Extending and evaluating a warfarin dosing algorithm that includes CYP4F2 and pooled rare variants of CYP2C9. *Pharmacogenetics and Genomics* 20, 7 (2010), 407-413.

EVENTS

CONFERENCES

ACM-SIAM Symposium on Discrete Mathematics (SODA14)
Hilton Portland & Executive Tower
Portland, OR
January 5-7, 2014
http://www.siam.org/meetings/alenex14

Ninth International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory, and Applications (VISIGRAPP 2014)
Sana Lisbon
Lisbon, Portugal
January 5-8, 2014
http://www.visigrapp.org

27th International Conference on VLSI Design & 13th International Conference on Embedded Systems (VLSI '14)
IIT Bombay, Victor Menezes Convention Center
Mumbai, India
January 5-9, 2014
http://www.vlsidesignconference.org

International Conference on Model-Driven Engineering and Software Development (MODELSWARD'14)
Sana Lisbon
Lisbon, Portugal
January 7-9, 2014
http://www.modelsward.org

International Conference on Physiological Computing Systems (PhyCS '14)
Sana Lisbon
Lisbon, Portugal
January 7-9, 2014
http://www.phycs.org

COMSNETS '14: International Conference on Communication Systems and Networks
The Chancery Pavilion
Bangalore, India
January 7-10, 2014
http://www.comsnets.org

The Eigth International Conference on Ubiquitous Information Management and Communication (ICUIMC '14)
Sofitel Angkor Phokeethra Golf and Spa Resort
Siem Reap, Cambodia
January 9-11, 2014
http://www.icuimc.org

2014 International Conference on Electronic Systems, Signal Processing and Computing Technologies (ICESC)
Shri Ramdeobaba College of Engineering and Management
Nagpur, India
January 9-11, 2014
http://www.rcoem-icesc.com

2014 IEEE International Conference on Consumer Electronics (ICCE)
Las Vegas Convention Center
Las Vegas, NV
January 10-13, 2014
www.icce.org

Innovations in Theoretical Computer Science (ITCS'14)
Princeton, NJ
January 12-14, 2014
http://www.wisdom.weizmann.ac.il/~naor/itcs_cfp.html

2014 Joint Mathematics Meeting
Baltimore Convention Center
Baltimore, MD
January 15-18, 2014

http://jointmathematicsmeetings.org/jmm

Cryptography and Security in Computing Systems
Vienna, Austria
January 20, 2014
http://www.cs2.deib.polimi.it

19th Asia and South Pacific Design Automation Conference (ASPDAC 2014)
International Convention & Exhibition Center
Suntec, Singapore
January 20-23, 2014
http://www.ece.nus.edu.sg/stfpage/elehy/aspdac2014

16th Australasian Computing Education Conference
Auckland City, New Zealand
January 20-23, 2014
http://elena.ait.ac.nz/homepages/ace2014

The Fourth International Workshop on Adaptive Self-tuning Computing Systems (ADAPT '14)
Vienna, Austria
January 22, 2014
http://www.adapt-workshop.org

Sixth Workshop on Rapid Simulation and Performance Evaluation: Methods and Tools (RAPIDO '14)
Vienna, Austria
January 22, 2014
http://www.hipeac.net/rapido/2014/index.html

The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '14)
US Grant San Diego Hotel
San Diego, CA
January 22-24, 2014
http://popl.mpi-sws.org/2014

The Eighth International Workshop on Variability Modelling of Software-intensive Systems (VAMOS 2014)
Nice Sophia Antipolis University
Sophia Antipolis, France
January 22-24, 2014
http://vamos2014.unice.fr

2014 Sixth International Conference on Knowledge and Smart Technology
Burapha University
Chonburi, Thailand
January 30-31, 2014
http://www.kst-thailand.org

Richard Tapia Celebration of Diversity in Computing Conference (TAPIA '14)
Grand Hyatt Seattle
Seattle, WA
February 5-8, 2014
http://www.tapiaconference.org

2014 International Conference on Information Networking (ICOIN)
Phuket, Thailand
February 10-12, 2014
www.icoin.org

ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '14)
The Peabody Orlando
Orlando, FL
February 15-19, 2014
https://sites.google.com/site/ppopp2014

Computer Supported Cooperative Work and Social Computing (CSCW 2014)
Baltimore Marriott Waterfront
Baltimore, MD
February 15-19, 2014
http://cscw.acm.org

Eighth International Conference on Tangible, Embedded, and Embodied Interaction (TEI'14)
Ludwig Maximilian University of Munich
Munich, Germany
February 16-19, 2014
http://www.tei-conf.org/14

2014 IEEE Sensors Applications Symposium
Rydges Lakeland Resort
Queenstown, New Zealand
February 18-20, 2014
http://sensorapps.org/

SIAM Conference on Parallel Processing for Scientific Computing (PP14)
Marriott Portland
Downtown Waterfront
Portland, OR
February 18-21, 2014
http://www.siam.org/meetings/pp14/

19th International Conference on Intelligent User Interfaces (IUI'14)
Haifa, Israel
February 24-27, 2014
http://iuiconf.org

2014 IEEE Fifth Latin American Symposium on Circuits and Systems (LASCAS)
Hotel Plaza San Francisco
Santiago, Chile
February 25-28, 2014
http://www.ieee-lascas.org

The 15th International Workshop on Mobile Computing Systems and Applications
Santa Barbara, CA
February 26-27, 2014
http://www.hotmobile.org/2014

The 2014 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA '1)
Monterey Conference Center
Monterey, CA
February 26 - 28, 2014
http://fpganetworks.org/FPGA2014

International Symposium on Engineering Secure Software and Systems
Munich, Germany
February 26-28, 2014
https://distrinet.cs.kuleuven.be/events/essos/2014

CONTESTS & EVENTS

Wearable Technology at CES
Today, the world of apparels, accessories, and technology are converging at a high velocity. To keep up with this pace the 2014 international CES is all set to feature a novel exhibit area called "FashionWare" focusing on the latest, cool innovations in wearable technology. So, if you have ever fantasized about a jacket that adjusts itself based on outside temperature or a solar-charging handbag make it a point to be at CES. For dates and even more keep an eye on http://www.cesweb.org/Home.aspx

Smart Lighting 2014
With the introduction of semiconductor based digital light sources such as LEDs and OLEDs, the lighting industry experienced a paradigm shift. Marriage of semiconductor and lighting technologies has opened up an array of new functionalities and a myriad of exciting new applications. Lighting has now become dynamic, adaptable, and interactive. Smart Lighting 2014 invites new players in this emerging market, technologists, and researchers to utilize this platform to discover their peers. For more information visit http://www.smartlighting.org/sl2014/

FEATURED EVENT

Photograph by Andrew Zarivny

The 15th International Workshop on Mobile Computing Systems and Applications
Santa Barbara, CA
February 26-27, 2014

Since the beginning of time, technology has been crafted with the purpose of bettering society. Whether it was the invention of the wheel, the invention of the personal computer, or the invention of the mobile phone, technology has revolutionized our everyday lives. This sort of technology and computing has transcended all limits and redefined what is possible in the world today, especially with the use of mobile technologies in everyday activities.

The International Workshop on Mobile Computing Systems and Applications is an event that focuses on mobile applications—one of the fastest growing areas in today's technology. This small and selective workshop focuses on mobile applications, environments, new technologies, and controversial directions. Perfect for research papers dealing with modern and untraditional computing, this conference will be hold for two days in the beautiful city of Santa Barbara.

For more information, please visit http://www.hotmobile.org/2014/.

—Rohit Goyal

ACRONYMS

AmI Ambient Intelligence: A vision on the future of consumer electronics, telecommunications, and computing that was originally developed in the late 1990s. It represents an electronic environment that is sensitive and responsive to the presence and the activity of people. Smart homes can be considered part of this vision.

AR Augmented Reality: A live view of a physical environment whose elements are "augmented" by computer-generated sensory input (sound, video, etc.). Typical devices using AR technologies are smartphones, tablets, and head-mounted displays.

BCI Brain Computer Interface : The direct communication path between the brain and an external device; also known as a mind-machine interface or direct neural interface.

EEG Electroencephalography: The recording of neural electrical activity using electrodes placed on the scalp. Not to be confused with electrooculography (EOG), which involves measuring retinal activity via electrodes.

OHMD Optical head-mounted display: A wearable display that has the capability of reflecting projected images and allowing users to see-through them. A practical and recent example of OHMD is Google Glass.

UbiComp Ubiquitous Computing: A computing concept, similar to AmI, where computing is made to appear everywhere and anywhere, in form of computers, smartphones, watches, etc.

GRANTS, SCHOLARSHIPS & FELLOWSHIPS

Department of Energy Computational Science Graduate Fellowship
Website: *https://www.krellinst.org/doecsgf/application/*
Deadline: January 7, 2014
Eligibility: US citizens or permanent residents in their first year of graduate work or senior year of undergraduate studies.
Benefits: $36,000 stipend, tuition and fees, and an allowance for a computer workstation
Explanation: This fellowship funds students in computational science who use high-performance computing to solve problems in science and engineering. This fellowship also includes professional development benefits with the U.S. Department of Energy.

Naval Research Enterprise Internship Program
Website: *http://nreip.asee.org/*
Deadline: January 6, 2014
Eligibility: U.S. citizens who have finished at least their sophomore year in college.
Benefits: $5,400 - $10,800 depending on level of study
Explanation: This program provides 10-week internships at a U.S. Navy laboratory to students.

Anita Borg Scholarship
Website: *http://www.google.com/anitaborg*
Deadline: Spring 2014
Eligibility: Female computer science or computer engineering students in their senior year of college or are graduate students
Benefits: $10,000 and attendance at Google Scholars Retreat
Explanation: Google offers this scholarship to support women in technology and encourage them to become leaders in the field.

GRC Graduate Fellowship Program
Website: *http://www.src.org/student-center/fellowship/#gfp*
Deadline: February 2014
Eligibility: U.S. citizens or permanent residents with at least two years remaining in a Ph.D. program researching areas relevant to microelectronics
Benefits: Full tuition and fees, with stipend
Explanation: Recipients will be matched with an industry advisor and will be given assistance finding a position at an SRC company, government agency, or academia. The award can be for up to five years.

POINTERS

WEARABLE COMPUTING RESOURCES

The economist Robert Solow was at least half right when he quipped, "You can see the computer age everywhere but in the productivity statistics." The advent of Google Glass hails an era where technology has become integral to human networking, communication, and maybe even productivity. We've provided a few resources for you to learn more about the emergence of wearable tech in modern life.
—*Ashok Rao*

READING LIST

Make: Wearable Electronics – Tools and Techniques for Prototyping Interactive Wearables
Kate Hartman, Maker Media, Inc. (2013)
Kate Hartman, an assistant professor at OCAD University, examines the hardware design process for totally new kind of technology.

"Our bodies are our primary interface for the world. Interactive systems designed to be worn can be intimate, upfront, and sometimes in your face (literally). Bringing wearable electronics from concept to prototype to product can be both inspiring and challenging. This book gives you what you need to start working with these new materials, tools, and techniques. It covers popular wearable products such as the Arduino Lilypad, Adafruit Flora, and the Fabrickit." (Publisher's description)

"We're Using a Ton of Mobile Data. With Google Glass, We're About to use a Whole Lot More."
By Brian Fung
Once Google Glass is available to consumers, the influx of users may crash our mobile networks. Fung investigates this potential burden on the nation's Internet infrastructure. *http://www.washingtonpost.com/blogs/the-switch/wp/2013/07/30/were-using-a-ton-of-mobile-data-with-google-glass-were-about-to-use-a-whole-lot-more/*

"Chinese Consumers Excited by Wearable Technology"
According to a recent survey of middle-income Chinese consumers, the demand for wearable technology is on the rise in China. As the industry tries to gain a stronger foothold in the U.S., Chinese consumers are aware of and excited about the technology. With health and fitness tracking devices garnering the most interest. *http://news.yahoo.com/chinese-consumers-excited-wearable-technology-170645670.html*

"Google Glass Privacy Concerns Persist in Congress"
By Charlie Osborne
Despite consumer interests in wearable technology, privacy is still an important concern for many. U.S. Rep. Joe Barton of Texas, is one of many, who has expressed disappointment with Google in its handling of privacy inquiries. *http://news.cnet.com/8301-1023_3-57591975-93/google-glass-privacy-concerns-persist-in-congress/*

WEBSITES

The Switch
The *Washington Post*'s business blog discusses the connection between technology and policy, producing ever relevant information for those interested in technology—wearable or not. As government policy becomes increasingly relevant for privacy and net neutrality in a new world, this is an excellent blog to follow. *http://www.washingtonpost.com/blogs/the-switch/*

Ecouterre
While this is a fashion site in spirit, many of their most interesting articles are about using technology to better everyday life. With catchy headlines such as "Tiny, Clothing-Embedded Cameras Could Help Dieters Track Calories" or the "Smartwatch [that] Turns Your Wrist Into a Phone," Ecouterre has an extensive archive of articles addressing wearable technology. *http://www.ecouterre.com/category/wearable-technology/*

Wearable Tech World
Featuring expert opinions, product reviews, and industry news, this should be your final destination for all things wearable. Not only a news source, there is also an annual conference, Wearable Tech Expo. The event strikes at the heart of wearable technology, featuring a variety of speakers from fashion, technology, and design. *http://www.wearabletechworld.comwearable-technology/*

BEMUSEMENT

Puzzles: Door #100

You have 100 doors in a row that are all initially closed. You make 100 passes by the doors starting with the first door every time. The first time through, you visit every door and toggle the door (if the door is closed, you open it; if it's open, you close it). The second time you only visit every second door (door #2, #4, #6, etc.). The third time, every third door (door #3, #6, #9, etc.), and so on, until you only visit the 100th door. What state are the doors in after the last pass? Which are open and which are closed?

Source: http://www.techinterview.org/post/526370758/100-doors-in-a-row

Final Score

Jennifer took a test that had 20 questions. The total grade was computed by awarding 10 points for each correct answer and deducting five points for each incorrect answer. Jennifer answered all 20 questions and received a score of 125. How many wrong answers did she have?

Source: Puzzle #3 at http://malini-math.blogspot.com/2009/08/simple-math-puzzles.html

Find the solution at: http://xrds.acm.org/bemusement/2013.cfm

Exercise vs. Time

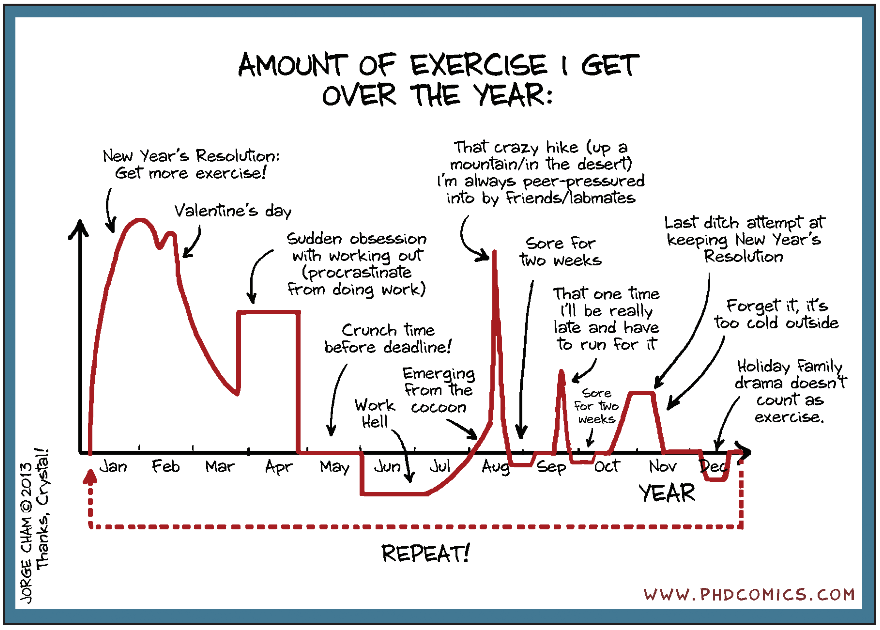

PhD Comics ©Jorge Cham

Anti Glass

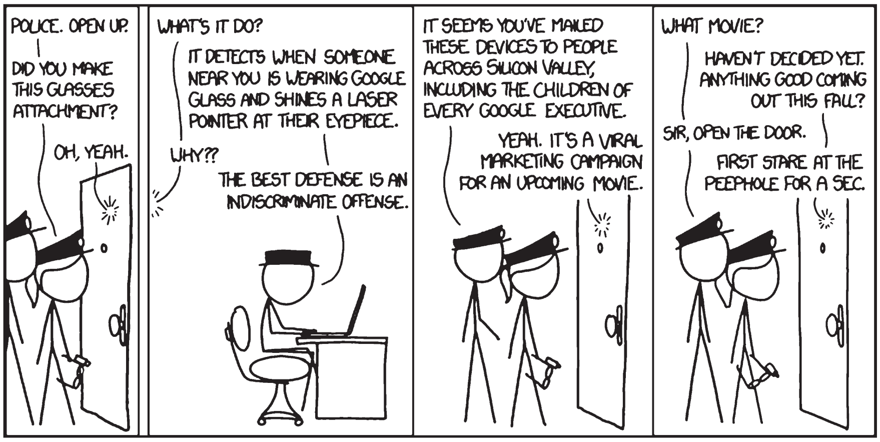

http://xkcd.com/1251/

SUBMIT A PUZZLE

Can you do better? Bemusements would like your puzzles and mathematical games (but not Sudoku). Contact **xrds@acm.org** to submit yours!

STUDENT MEMBERSHIP APPLICATION

Join ACM online: www.acm.org/joinacm

CODE: CRSRDS

Name — *Please print clearly*

Address

City — State/Province — Postal code/Zip

Country — E-mail address

Area code & Daytime phone — Mobile phone — Member number, *if applicable*

MEMBERSHIP BENEFITS AND OPTIONS

- Free software and courseware through the ACM Academic Initiative
- Free e-mentoring services from MentorNet®
- Electronic subscriptions to *Communications of the ACM* and *XRDS: Crossroads* magazines
- Online courses, online books and videos
- ACM's *CareerNews* (twice monthly)
- ACM e-news digest *TechNews* (thrice weekly)
- ACM online newsletter *MemberNet* (monthly)
- *Student Quick Takes*, ACM student e-newsletter (quarterly)
- Free "acm.org" email forwarding address plus filtering through Postini
- Option to subscribe to the full ACM Digital Library
- Discounts on ACM publications and conferences, valuable products and services, and more

PLEASE CHOOSE ONE:

❏ **Student Membership: $19 (USD)**
❏ **Student Membership PLUS Digital Library: $42 (USD)**
❏ **Student Membership PLUS Print *CACM* Magazine: $42 (USD)**
❏ **Student Membership w/Digital Library PLUS Print *CACM* Magazine: $62 (USD)**

PUBLICATIONS

Please check

Check the appropriate box and calculate amount due on reverse.	Issues per year	Code	Member Rate	Air Rate*
• ACM Inroads	4	178	$16 ❐	$58 ❐
• Communications of the ACM	12	101	$25 ❐	$58 ❐
• Computers in Entertainment (online only)	4	247	$43 ❐	N/A
Computing Reviews	12	104	$55 ❐	$39 ❐
• Computing Surveys	4	103	$36 ❐	$32 ❐
Evolutionary Computation (MIT Press)	4	177	$32 ❐	$30 ❐
• interactions, new visions of human-computer interaction (included in SIGCHI membership)	6	123	$22 ❐	$35 ❐
• Int'l Journal of Network Management (online only) (Wiley)	6	136	$92 ❐	$30 ❐
Int'l Journal on Very Large Databases	4	148	$83 ❐	$30 ❐
• Journal of Educational Resources in Computing (see TOCE)	N/A	N/A	N/A	N/A
• Journal of Experimental Algorithmics (online only)	12	129	$31 ❐	N/A
• Journal of Personal and Ubiquitous Computing	6	144	$65 ❐	$30 ❐
• Journal of the ACM	6	102	$55 ❐	$58 ❐
• Journal on Computing and Cultural Heritage	4	173	$49 ❐	$23 ❐
• Journal on Data and Information Quality	4	171	$49 ❐	$25 ❐
• Journal on Emerging Technologies in Computing Systems	4	154	$42 ❐	$23 ❐
• Linux Journal (SSC)	12	137	$31 ❐	$33 ❐
• Mobile Networks and Applications	6	130	$72 ❐	$29 ❐
• Wireless Networks	4	125	$72 ❐	$29 ❐
• XRDS (included with membership)	4	XRoads	$39 ❐	N/A
Transactions on:				
• Accessible Computing	4	174	$49 ❐	$24 ❐
• Algorithms	4	151	$52 ❐	$23 ❐
• Applied Perception	4	145	$43 ❐	$23 ❐
• Architecture & Code Optimization	4	146	$43 ❐	$23 ❐
• Asian Language Information Processing	4	138	$39 ❐	$23 ❐
• Autonomous and Adaptive Systems	4	158	$41 ❐	$23 ❐
• Computational Biology and Bioinformatics	4	149	$20 ❐	$49 ❐
• Computer-Human Interaction	4	119	$43 ❐	$25 ❐
• Computational Logic	4	135	$44 ❐	$25 ❐
• Computation Theory	8	176	$49 ❐	$32 ❐
• Computer Systems	4	114	$47 ❐	$25 ❐
• Computing Education (formerly JERIC)		277	$25 ❐	N/A
• Database Systems	4	109	$46 ❐	$25 ❐
• Design Automation of Electronic Systems	4	128	$43 ❐	$25 ❐
• Economics and Computation	4	192	$49 ❐	$23 ❐
• Embedded Computing Systems	4	142	$44 ❐	$23 ❐
• Graphics	4	112	$51 ❐	$25 ❐
• Information and System Security	4	134	$44 ❐	$23 ❐
• Information Systems	4	113	$47 ❐	$25 ❐
• Intelligent Systems and Technology	4	179	$46 ❐	$72 ❐
• Interactive Intelligent Systems	4	191	$49 ❐	$76 ❐
• Internet Technology	4	140	$42 ❐	$23 ❐
• Knowledge Discovery From Data	4	170	$50 ❐	$23 ❐
• Management Information Systems	4	190	$47 ❐	$22 ❐
• Mathematical Software	4	108	$47 ❐	$25 ❐
• Modeling and Computer Simulation	4	116	$51 ❐	$25 ❐
• Multimedia Computing, Communications, and Applications	4	156	$42 ❐	$23 ❐
• Networking	6	118	$29 ❐	$52 ❐
• Programming Languages & Systems	6	110	$59 ❐	$32 ❐
• Reconfigurable Technology & Systems	4	172	$49 ❐	$24 ❐
• Sensor Networks	4	155	$42 ❐	$23 ❐
• Software Engineering and Methodology	4	115	$43 ❐	$25 ❐
• Speech and Language Processing (online only)	4	253	$33 ❐	N/A
• Storage	4	157	$42 ❐	$23 ❐
• Web	4	159	$41 ❐	$23 ❐

Marked • are available in the ACM Digital Library
** Check here to have publications delivered via Expedited Air Service. For residents outside North America only.*

PUBLICATION SUBTOTAL: ______

INSTRUCTIONS

Carefully complete this application and return with payment by mail or fax to ACM. You must be a **full-time** student to qualify for student rates.

CONTACT ACM

phone: 800-342-6626 (US & Canada)
+1-212-626-0500 (Global)
hours: 8:30am–4:30pm US Eastern Time
fax: +1-212-944-1318
email: acmhelp@acm.org
mail: Association for Computing Machinery, Inc.
General Post Office
P.O. Box 30777
New York, NY 10087-0777

For immediate processing, FAX this application to +1-212-944-1318.

PAYMENT INFORMATION

Payment must accompany application

Member dues ($19, $42, or $62) $ ______

To have *Communications of the ACM* sent to you via Expedited Air Service, add $58 here (for residents outside of North America only). $ ______

Publications $ ______

Total amount due $ ______

Check or money order (make payable to ACM, Inc. in U.S. dollars or equivalent in foreign currency)

❏ **Visa/Mastercard** ❏ **American Express**

Card number — Exp. date

Signature

Member dues, subscriptions, and optional contributions are tax deductible under certain circumstances. Please consult with your tax advisor.

EDUCATION

Name of School

Please check one: ❐ **High School** (Pre-college, Secondary School) **College:** ❐ Freshman/1st yr. ❐ Sophomore/2nd yr. ❐ Junior/3rd yr. ❐ Senior/4th yr. **Graduate Student:** ❐ Masters Program ❐ Doctorate Program ❐ Postdoctoral Program ❐ **Non-Traditional Student**

Major — Expected mo./yr. of grad.

Age Range: ❐ 17 & under ❐ 18-21 ❐ 22-25 ❐ 26-30 ❐ 31-35 ❐ 36-40 ❐ 41-45 ❐ 46-50 ❐ 51-55 ❐ 56-59 ❐ 60+

Do you belong to an ACM Student Chapter? ❐ Yes ❐ No

I attest that the information given is correct and that I will abide by the ACM Code of Ethics. I understand that my membership is non transferable.

Signature