

Tomas Petricek – Dissertation Preface

<http://tomasp.net> | tomas@tomasp.net

Most of my research to date has been focused on the subject of programming languages and systems. The topic can be studied from multiple perspectives and in multiple contexts and my work explored multiple such research directions, each spanning multiple individual projects and papers.

I studied the theory of (i) *context-aware programming* (in my PhD thesis), worked on (ii) *functional programming* and the F# language (in parallel to my PhD studies), programming in the context of (iii) *data exploration* (during my two Post-doctoral stays) and researched (iv) *interactive programming systems* (since joining University of Kent). I also devoted a part of my time to (v) *interdisciplinary research*, focusing on the history and philosophy of programming.

In each of those areas, my work has had (or is showing early signs of having) a lasting academic and industrial impact. The presented annotated selection of the five most important works serves as the evidence for this fact. It includes five significant academic works, each representing one of the five aforementioned research areas. The five selected works have been widely cited, received significant academic awards, inspired international grant applications, influenced main-stream programming languages or have been subject to invited keynotes.

My Habilitation dissertation *Simple programming tools for data exploration* is centred on one of the five research directions outlined above, choosing depth over breadth. It presents my work on data exploration, which is an independent research direction that I pursued following my PhD and for which I secured funding from multiple competitive sources. The research employed theoretical programming language knowledge in the practical area of working with data. As such, the work is of a more applied nature than some of my earlier research.

The dissertation provides a comprehensive summary of my work on data exploration, including key theoretical publications, publications accompanying influential software artifacts, as well as publications that complement the core theory with the practically necessary applied developments. The aim of the dissertation is to show the work in context and highlight connections that may not be apparent when reading individual papers. I also use the longer format of a dissertation to highlight interesting work that has been perhaps less widely recognized.

The key theoretical results in the dissertation rethink established programming language notions or use them in a new way. These have been published in top conferences including CORE A* (Chapters 6 and 13) and CORE A (Chapter 7), indexed journals (Chapters 11, 12) or received best paper awards (Chapter 6 and 8). The theoretical results have been accompanied by multiple open-source software packages, some of which have become widely adopted by industry.

The focus on the problem of data exploration in the dissertation serves not only to provide a comprehensive overview of my recent work, but it also provides the background, and a fruitful motivating example, for my current research direction.

The work presented in the dissertation views data exploration as a programming language problem and uses a range of established programming language techniques to improve data science tools. However, this perspective also raises new basic questions about programming language research methodologies. It shows that we need to consider not only the languages, but also the stateful interactive systems in which they are embedded. Generalizing from *programming languages* to *programming systems* is an intriguing research direction that I intend to continue exploring in the future.